

REPORT NUMBER 23

**LINEAR TIME-FREQUENCY REPRESENTATIONS FOR TRANSIENT
SIGNAL DETECTION AND CLASSIFICATION**

N. LEE AND S.C. SCHWARTZ

INFORMATION SCIENCES AND SYSTEMS LABORATORY

**Department of Electrical Engineering
Princeton University
Princeton, New Jersey 08544**

June 1995

DTIC QUALITY INSPECTED 4

Prepared for

**OFFICE OF NAVAL RESEARCH (Code 313SP)
Surveillance, Communications and Electronic Combat Division
Arlington, Virginia 22217
under Contract N00014-91-J-1144**

S.C. Schwartz, Principal Investigator

Approved for public release; distribution unlimited

19960606 046

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 23	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Linear Time-Frequency Representations for Transient Signal Detection and Classification		5. TYPE OF REPORT & PERIOD COVERED Technical Report Sept. 1992 - December 1994
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Nigel Lee and Stuart Schwartz		8. CONTRACT OR GRANT NUMBER(s) N00014-91-J-1144
9. PERFORMING ORGANIZATION NAME AND ADDRESS Information Sciences & Systems Laboratory Dept. of Electrical Engineering Princeton University, Princeton, NJ 08544		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research (Code 313SP) Department of the Navy Arlington, VA 22217		12. REPORT DATE June 1995
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		13. NUMBER OF PAGES 152
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES Also submitted as the Ph.D. thesis of Nigel Lee to the EE Department, Princeton University, Princeton, NJ 08544, December 1994		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Gabor representation oversampling transient detection transient classification		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This dissertation examines the use of linear time-frequency representations in the detection and classification of transient signals. In particular, the Gabor transform and short-time Fourier transform (STFT) are shown to be effective tools in detecting and classifying signals that are accurately modeled by linear subspaces. Continued		

Continued

Transient signal detection is studied first, within a framework that expresses transient signals as linear combinations of time-frequency-shifted, one-sided exponential window functions. For the case where signal components have known locations in the time-frequency plane, it is shown that a generalized likelihood ratio test (GLRT) detector based on the oversampled Gabor transform is more robust to mismatch than GLRT detectors based on the critically sampled Gabor transform and critically sampled STFT. For the case where signal component locations are not precisely known, it is shown that, for a given transform, a GLRT detector which does not make assumptions about component location information is more robust to component location mismatch than a GLRT detector which does make those assumptions.

When the oversampled Gabor transform is used for data reduction, one of its main drawbacks is its lack of stability: small variations in a signal can cause large variations in the magnitudes of the Gabor coefficients. Thus, several modifications designed to improve the stability of the oversampled Gabor transform have been widely used, and the transform is stable in this form. However, it is shown here that there are several serious problems with the expanded form of the oversampled Gabor transform that make it unsuitable for use in transient signal detection.

Next, the oversampled Gabor transform is applied to a simple transient signal classification problem. Two GLRT-based classifiers based on the oversampled Gabor transform are shown to be largely ineffective. Because of this, a classifier based on the adaptive Gabor transform is defined, and it is shown that this classifier provides performance far superior to that of the GLRT-based classifiers. Finally, the optimal detector-classifier combination for the problems studied in this dissertation is proposed.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Overview	2
2	Preliminaries	4
2.1	Overview of Time-Frequency Representations	4
2.1.1	Quadratic TFRs: properties and examples	5
2.1.2	Linear TFRs: properties and examples	9
2.2	Overview of the Gabor Transform	12
2.2.1	Continuous-time Gabor transform	12
2.2.2	Discrete-time Gabor transform for finite-length signals	16
3	The Gabor Transform In Transient Signal Detection	18
3.1	Introduction	18
3.2	A Framework for Transient Signal Detection	19
3.3	Calculating the Gabor Transform Within the PF Framework	22
3.3.1	Discretized biorthogonal function method	22
3.3.2	Least squares method	23
3.3.3	General transforms within the PF framework	24

3.4	Transient Detection Within the PF Framework	24
3.4.1	Hypothesis test and GLRT statistic	24
3.4.2	Analysis of GLRT statistic within PF framework	27
3.5	PF Analysis Results and Suggested Improvement: Oversampling . . .	29
3.5.1	Porat-Friedlander results	29
3.5.2	Oversampling to improve GLRT robustness in the presence of location mismatch	32
3.6	Comparison of Detector Performances	36
3.6.1	Modified framework	37
3.6.2	Location mismatch	39
3.6.3	Shape mismatch	43
3.6.4	Effects of varying the number K of signal components	44
3.6.5	Oversampling the other linear transforms	45
3.7	Detection of Unknown-Location Signal Components in the PF Frame- work	49
3.8	Analysis: UL Detector vs. KL Detector	52
3.8.1	Simulation examples	53
3.8.2	Comparison: UL detector and energy detector	62
3.9	Scharf-Friedlander Framework: Matched Subspace Detectors	63
3.10	Conclusions	65
3.A	Appendix: Restrictions on ν_{min} in Section 3.4	65
4	Improving the Oversampled Gabor Transform	68
4.1	Introduction	68
4.2	Stability Analysis of TFRs in the PF Framework	69
4.3	Methods of Improving Coefficient Stability in the OS-LSM-GT	75
4.3.1	Rank reduction	75
4.3.2	Regularization	76
4.3.3	Principal components method	80

4.3.4	LSM solution with energy constraints	83
4.3.5	LSM solution with transform "whitening"	84
4.4	Performance Analysis of Modifications to the OS-LSM-GT	87
4.4.1	Performance measures	87
4.4.2	Analysis results	88
4.5	Wexler-Raz (WR) Framework for Discrete-Time Gabor Transforms .	93
4.6	Calculation of the Oversampled Gabor Transform in the WR Framework	95
4.6.1	Periodized biorthogonal function method	96
4.6.2	Direct solutions	98
4.6.3	Examples	99
4.7	Analysis of the Oversampled Gabor Transform in the WR Framework	104
4.7.1	Stability analysis	104
4.7.2	Accuracy analysis	104
4.8	Conclusions	108
5	The Gabor Transform In Transient Signal Classification	110
5.1	Introduction	110
5.2	Framework for Transient Signal Classification	111
5.3	GLRT-Based Classifiers for Transient Signal Classification	112
5.3.1	GLRT classifier	112
5.3.2	Modified-GLRT classifier	118
5.4	Adaptive Gabor Transform (AGT)	120
5.4.1	Calculating the AGT	120
5.4.2	Examples and performance measures for AGT	124
5.5	Simulation Results	130
5.5.1	AGT classifier	130
5.5.2	Monte Carlo simulations	131
5.5.3	Optimal detector-classifier combination	134

5.6 Conclusions	135
---------------------------	-----

6 Summary and Future Work	137
----------------------------------	------------

Introduction

1.1 Motivation

This dissertation examines the use of linear time-frequency representations in the detection and classification of transient signals. Time-frequency representations (TFRs) combine the traditional elements of time-domain analysis and frequency-domain analysis to form a joint time-frequency characterization of a given signal. For *stationary* signals whose spectral contents remain essentially constant over time, standard Fourier analysis is sufficient and the additional information provided by TFRs is unnecessary. However, many important signal processing applications today — in areas as diverse as medicine, geophysics, defense, and music — involve *nonstationary* signals, signals whose spectral contents vary with time. In analyzing these types of signals, it is necessary to determine *which* frequency components occurred *when*, and this is what TFRs are designed to do.

Transient signals are a commonly encountered class of nonstationary signals. A *transient signal* is simply defined as a nonstationary signal whose duration is short compared to the overall observation time [FP89]. Optimal detection of a known transient signal is accomplished by means of the familiar matched filter. When the transient signal is not completely known (which is almost always the case), it is necessary to resort to maximum likelihood or generalized likelihood techniques. In this case, it is advantageous to transform the received data before the detection process begins so that the detector is based not on a very large number of data samples but on a smaller number of transform coefficients. This makes the detection process both simpler and faster. Clearly, a given transform will only be effective in a transient signal detection scheme if it represents the information content of the signal *accurately* as well as simply. For this reason, the Fourier transform is *not* effective

when applied to transient signal detection problems, because it cannot distinguish signal components that have the same frequency but occur at different times.

The greater part of this dissertation focuses on the application of a particular linear TFR, the Gabor transform, to transient signal detection and classification schemes. By nature, the Gabor transform is designed to represent the information content of a given signal with a relatively small number of coefficients. Furthermore, when the shapes of the desired signal components are known, or known up to a set of parameters, the Gabor transform coefficients represent the information content of the signal quite accurately as well. Thus, the Gabor transform is particularly well-suited for use in transient signal detection and classification problems.

1.2 Overview

The body of this dissertation is organized into five chapters as summarized below.

Chapter 2 gives some necessary background material on TFRs in general and on the Gabor transform in particular. In the general overview of TFRs, much attention is devoted to the distinction between *linear* and *quadratic* TFRs, and it is explained why quadratic TFRs such as the Wigner distribution are not suitable for use in transient detection problems. In the discussion on the Gabor transform, the emphasis is placed on the various means of calculating the Gabor transform coefficients and on the advantages and drawbacks to each method. Also, the distinction is made between Gabor transform coefficients that are *critically sampled* and *oversampled*.

Chapter 3 applies the Gabor transform to a class of transient signal detection problems. The first part of the chapter makes use of a framework first developed by Friedlander and Porat [FP89, PF92] which assumes that the desired signal components have known locations in the time-frequency plane and known waveform shapes, but unknown magnitudes. Within this framework, the performance of a GLRT¹ detector based on the oversampled Gabor transform is compared to the performances of four detectors based on critically sampled linear transforms. The main performance criterion is *robustness* to signal mismatch, i.e., the ability of a detector to maintain its effectiveness when the actual parameter values of the signal components vary from the assumed parameter values. It is shown that the oversampled Gabor transform detector is most robust among all the detectors to mismatches in both location and waveform shape.

In the second part of Chapter 3, a GLRT detector based on the oversampled Gabor transform is derived for the case where the locations of the desired signal

¹Generalized likelihood ratio test

components are *not* assumed to be known. The performance of this unknown-location (UL) detector is then compared to the performance of the corresponding known-location (KL) GLRT detector derived in the first part of the chapter. It is shown that, as expected, the UL detector is more robust to location mismatch than the KL detector.

Chapter 4 explores various methods for improving upon the main weaknesses of the oversampled Gabor transform, its lack of coefficient stability (meaning that the coefficient magnitudes show large variations) in the presence of noise or signal mismatch. First, an explanation is given as to why the GLRT detector based on the oversampled Gabor transform proved to be robust to signal mismatch in Chapter 3 even though the oversampled Gabor transform coefficients themselves are unstable in the presence of mismatch. Then, several methods of counteracting this instability are studied, including rank reduction (via singular value decomposition), regularization, solutions with energy constraints, and the principal components method. The last part of Chapter 4 examines the calculation of the oversampled Gabor transform under a different set of assumptions, first proposed by Wexler and Raz [WR90]. It is shown that the Wexler-Raz assumptions greatly improve the coefficient stability of the oversampled Gabor transform but greatly increase its computational complexity as well. Among other reasons, the latter makes the Wexler-Raz version of the oversampled Gabor transform unsuitable for use in transient signal detection problems.

Chapter 5 investigates a slightly different problem, that of transient signal classification. Here, it is assumed that a transient signal has been detected, but the problem is to determine what type of signal it is. A simple classification problem is proposed within the Friedlander-Porat framework, and a GLRT detector based on the oversampled Gabor transform is derived for this problem. However, it is shown that the GLRT detector is essentially useless for this problem, and an alternate detector (still based on the oversampled Gabor transform) is proposed, with much better results. Finally, a recently proposed method of calculating the Gabor transform coefficients adaptively, called *matching pursuit*, is applied to the classification problem, and it is shown through Monte Carlo simulations that a detector based on this method far outperforms either of the first two detectors. The implications of this are examined and the optimal detector-classifier combination is proposed.

Finally, Chapter 6 summarizes the conclusions drawn from Chapters 3 through 5 and suggests possible further avenues of research. The latter includes in particular a discussion of how to extend the results of this dissertation involving the Gabor transform to research problems that involve another recently popular linear TFR, the wavelet transform.

Chapter 2

Preliminaries

This chapter is intended to review the background material necessary to understand the problems and results of Chapters 3 through 5. The first half of the chapter gives a general overview of time-frequency representations (TFRs), while the second half presents an overview of the Gabor transform, the linear TFR that will be the primary focus of this dissertation. The respective overviews are by no means comprehensive; instead, they focus on material relevant to the problems examined in the subsequent chapters.

2.1 Overview of Time-Frequency Representations

Hlawatsch and Boudreaux-Bartels aptly describe the general concept and purpose of time-frequency representations (TFRs):

Time-frequency representations (TFRs) of signals map a one-dimensional signal of time, $x(t)$, into a two-dimensional function of time and frequency, $T_x(t, f)$. Most TFRs are “time-varying spectral representations” which are similar conceptually to a musical score with time running along one axis and frequency along the other axis. The values of the TFR surface above the time-frequency plane give an indication as to which spectral components are present at which times. [HBB92]

Two main classes of TFRs have emerged in the literature: quadratic TFRs and linear TFRs. The differences between the two classes begin with their definitions. A *linear* TFR $L_x(t, f)$ is calculated as a function of the first-order term $x(t)$ (the signal

itself), and it obeys the superposition principle for a linear combination of signals $c_1x_1(t) + c_2x_2(t)$:

$$L_{c_1x_1+c_2x_2}(t, f) = c_1L_{x_1}(t, f) + c_2L_{x_2}(t, f). \quad (2.1)$$

A *quadratic* TFR $Q_x(t, f)$ is calculated as a function of the second-order term $x(t)x^*(t - \tau)$, $\tau \in \mathcal{R}$, and it does *not* obey the superposition principle for a linear combination $c_1x_1(t) + c_2x_2(t)$:

$$Q_{c_1x_1+c_2x_2}(t, f) = |c_1|^2Q_{x_1}(t, f) + |c_2|^2Q_{x_2}(t, f) + c_1c_2^*Q_{x_1x_2}(t, f) + c_2c_1^*Q_{x_2x_1}(t, f), \quad (2.2)$$

where the latter two “crossterms” define $Q(t, f)$ as a function of a cross-correlation $x_1(t)x_2^*(t - \tau)$ instead of an energy¹ $x(t)x^*(t - \tau)$. This fundamental difference in definitions between quadratic and linear TFRs results in other differences that will become apparent below as specific examples of each class are discussed.

2.1.1 Quadratic TFRs: properties and examples

One reason for defining TFRs as quadratic functions of the signal (instead of linear) is that one would like a TFR to represent the *energy distribution* of the signal in the time-frequency plane. As such, the two-dimensional TFR should be related to the one-dimensional energy distributions, the *instantaneous power* $p_x(t) = |x(t)|^2$ and the *spectral energy density* $P_x(f) = |X(f)|^2$ (where $X(f)$ is the Fourier transform of $x(t)$). Ideally, one should be able to obtain these “marginal” densities of one variable by integrating the TFR over the other variable²:

$$\int_f Q_x(t, f) df = p_x(t) \quad (2.3)$$

$$\int_t Q_x(t, f) dt = P_x(f) \quad (2.4)$$

$$\int_t \int_f Q_x(t, f) dt df = E, \quad (2.5)$$

where $E = \int_t p_x(t) dt = \int_f P_x(f) df$ is the *energy* of the signal. A TFR that satisfies the above three relations is said to “satisfy the marginals” [Coh89].

The **Wigner distribution** (WD) is a widely-used quadratic TFR that satisfies the marginals and has numerous other desirable properties [CM80, HBB92, Coh89].

¹Here, the terms *correlation* and *energy* are used loosely: one needs to integrate these terms over time t to obtain the quantities normally defined as “correlation” and “energy.”

²Unless otherwise noted, all integrals and summations are assumed to range from $-\infty$ to ∞ .

The Wigner distribution $W_x(t, f)$ of a signal $x(t)$ is defined as

$$WD_x(t, f) = \int_{\tau} x(t + \frac{\tau}{2}) x^*(t - \frac{\tau}{2}) e^{-j2\pi f\tau} d\tau. \quad (2.6)$$

The “cross-Wigner” distribution $W_{x,y}(t, f)$ between two signals $x(t)$ and $y(t)$ is defined similarly as

$$WD_{xy}(t, f) = \int_{\tau} x(t + \frac{\tau}{2}) y^*(t - \frac{\tau}{2}) e^{-j2\pi f\tau} d\tau. \quad (2.7)$$

Using the relation in equation (2.2), it is easily shown that

$$WD_{x+y}(t, f) = WD_x(t, f) + WD_y(t, f) + 2\text{Re}\{WD_{xy}(t, f)\},$$

where the last term is the crossterm and the first two terms in the sum are called *autoterms*.

Among the many favorable properties of the Wigner distribution is that it provides maximal autoterm *concentration* (i.e., maximal time-frequency localization of the autoterms) among those quadratic TFRs that satisfy the marginals [JP92]. Unfortunately, the crossterms in the Wigner distribution are substantial and almost always interfere with the interpretation of the Wigner distribution; in fact, crossterms are sometimes called “interference terms.” To illustrate the effects of crossterms, consider Figure 2.1, which shows the Wigner distribution of the sum of two sinusoids, $e^{j2\pi t} + e^{j18\pi t}$. Note that, in addition to the well-concentrated autoterms at $f = 1\text{Hz}$ and $f = 9\text{Hz}$, there is a substantial crossterm at $f = 4.5\text{Hz}$. In a simple example such as this, one can easily distinguish the autoterms and the crossterm, but the effects of crossterms are much worse when the analysis involves more intricate signals or noise.

In fact, the Wigner distribution of any multicomponent signal³ contains noticeable crossterms. Consequently, there are several related quadratic TFRs that attempt to smooth the crossterms of the Wigner distribution while maintaining a good number of the WD’s favorable properties. These include:

- the smoothed pseudo-Wigner distribution [HBB92]
- the Choi-Williams distribution [CW89]
- the reduced interference distributions [JW92]

³Here, a *multicomponent* signal is loosely defined as any signal $x(t)$ that may be expressed as a sum of two signals $x_1(t) + x_2(t)$, each obeying a separate “time-frequency law.” For a more precise definition, see [Coh92].




Figure 2.1: Wigner distribution of $e^{j2\pi t} + e^{j18\pi t}$. Note the presence of crossterms at $f = 4.5\text{Hz}$. The amplitude values should be ignored, as they involve finite approximations of delta functions. Example first used in [CW89].

- the cone-kernel distribution [ZAM90]
- the spectrogram [Alt80].

All of the above TFRs perform some type of smoothing on the WD, but none of them completely eliminate crossterms for every type of signal. In addition, studies have shown that there is a tradeoff between crossterm suppression and autoterm concentration [HBB92, JP92]. In other words, those quadratic TFRs that smooth crossterms well also tend to have poor concentration of autoterms. The Wigner distribution is at one extreme of the tradeoff, providing maximal autoterm concentration but no suppression of crossterms [JP92]. At the other extreme of the tradeoff is the spectrogram, which provides substantial smoothing of crossterms but poor autoterm concentration [KBB92]. The spectrogram is the squared magnitude of the short-time Fourier transform, a linear TFR discussed in the next subsection.

Since transient signal detection problems often involve multicomponent signals, the characterization of transient signals by means of quadratic TFRs is often problematic because of the presence of crossterms. Any quadratic TFR used to detect multicomponent transient signals must provide good suppression of crossterms. Thus, the spectrogram is as suitable as any other quadratic TFR for use in transient signal detection, while the Wigner distribution by itself (i.e., with no smoothing) is unsuitable. This will be discussed further in the discussion on linear TFRs below.

Finally, it should be noted that Cohen has formulated a unified approach to the class of quadratic, shift-invariant TFRs [Coh66, Coh89]. A TFR $T(t, f)$ is defined as *shift-invariant* if

$$y(t) = x(t - t_0)e^{2\pi f_0 t} \Rightarrow T_y(t, f) = T_x(t - t_0, f - f_0).$$

The class of quadratic, shift-invariant TFRs is sometimes known as *Cohen's class*, and any member of Cohen's class may be derived from the Wigner distribution by means of a time-frequency convolution:

$$T_x(t, f) = \int_{t'} \int_{f'} \Psi_T(t - t', f - f') W D_x(t', f') dt' df', \quad (2.8)$$

where the kernel function $\Psi_T(t, f)$ corresponding to $T(t, f)$ acts as a 2-D filter. All of the quadratic TFRs in the list above may be derived from the WD by means of equation (2.8), using the appropriate kernel function.

2.1.2 Linear TFRs: properties and examples

As stated previously, a linear TFR is defined as a TFR that obeys the superposition principle (see equation (2.1)). There are three linear TFRs that are widely used: the short-time Fourier transform (STFT), the wavelet transform (WT) and the Gabor transform. The Gabor transform will be discussed in a separate overview in the next section of this chapter, so the focus here will be on the STFT and on the WT.

The short-time Fourier transform

The STFT of a signal $x(t)$ is defined as

$$STFT_x(t, f) = \int_{t'} x(t') \gamma^*(t' - t) e^{-j2\pi f t'} dt'. \quad (2.9)$$

It is often thought of as a “windowed Fourier transform,” with the function $\gamma^*(t')$ serving as a sliding window. The STFT suffers from an inherent resolution tradeoff: good time resolution of the STFT requires $\gamma^*(t')$ to be narrow in time (i.e., a “short” window), while good frequency resolution of the STFT requires $\gamma^*(t')$ to be narrow in frequency (or narrowband, implying a “long” window in time). Unfortunately, the Heisenberg uncertainty principle⁴ [CM84, Coh89, LPA92] of Fourier analysis prohibits the existence of a window that has both good time resolution and good frequency resolution – hence the resolution tradeoff.

The synthesis of a signal $x(t)$ from its STFT is given by

$$x(t) = \int_{t'} \int_{f'} STFT_x(t', f') g(t - t') e^{j2\pi f' t} dt' df', \quad (2.10)$$

where $g(t)$ is known as the *synthesis window*, as opposed to the *analysis window* $\gamma^*(t')$ above. The synthesis and analysis windows must satisfy the fairly unrestrictive condition $\int_t g(t) \gamma^*(t) dt = 1$ in order for (2.9) and (2.10) to be consistent. It is common to choose the synthesis window equal to the analysis window, with an appropriate normalization.

Because it is linear, the STFT does not suffer from the presence of crossterms. However, the STFT is also complex-valued in general, making it difficult to use the STFT directly for transient signal detection problems unless one is prepared to

⁴This is indeed the same Heisenberg uncertainty principle that arises in quantum mechanics, involving position and momentum instead of time and frequency. As Cohen notes, there is at least a “partial formal mathematical correspondence between quantum mechanics and signal analysis.” See [Coh89] for details.

analyze complex hypothesis tests. More often, the complex-valued STFT is made real by taking its squared magnitude, known as the **spectrogram**:

$$SPEC_x(t, f) = |STFT_x(t, f)|^2. \quad (2.11)$$

The spectrogram is actually a quadratic TFR that can be derived from the Wigner distribution using equation (2.8) and an appropriate kernel, and it *does* contain crossterms. However, as noted earlier, the spectrogram provides substantial cross-term suppression (at the expense of autoterm resolution). In fact, Kadambe and Boudreaux-Bartels note that two signals whose spectrograms do not overlap in the time-frequency plane will have crossterms that are essentially zero [KBB92]. As will be detailed in subsequent chapters, it is usually assumed that transient signals are “far apart” in the time-frequency plane (far apart in time, at least), so detectors based on the spectrogram should not suffer from the presence of crossterms in transient signal detection applications.

The wavelet transform

The **wavelet transform** (WT) of a signal $x(t)$ is defined as

$$WT_x(t, f) = \sqrt{\left| \frac{f}{f_0} \right|} \int_{t'} x(t') \gamma^* \left(\frac{f}{f_0} (t' - t) \right) dt', \quad (2.12)$$

where the *analyzing wavelet* $\gamma^*(t')$ corresponds to the window function of the STFT and f_0 represents the center frequency of the analyzing wavelet. The differences between the STFT and the WT can be seen more easily if one treats $\frac{f_0}{f}$ in (2.12) above as a single quantity $a = \frac{f}{f_0}$, termed the *scale factor*. The WT of $x(t)$ is then rewritten as

$$WT_x(t, a) = \sqrt{\left| \frac{1}{a} \right|} \int_{t'} x(t') \gamma^* \left(\frac{1}{a} (t' - t) \right) dt'. \quad (2.13)$$

Actually, the WT is more appropriately referred to as a *time-scale* representation, as the frequency $f = af_0$ in the WT acts as part of the scaling process [RV91].

Close inspection of equation (2.9) reveals that the STFT of $x(t)$ at a particular time-frequency coordinate (t, f) is obtained by the following procedure:

1. Shift the window function $\gamma(t')$ by time t .
2. Modulate the time-shifted window function by $e^{j2\pi ft'}$.

3. Take the inner product⁵ of the signal with the time-shifted, frequency-modulated window function: $STFT_x(t, f) = \langle x, \gamma_{(t,f)} \rangle$, where $\gamma_{(t,f)}(t') = \gamma(t' - t)e^{j2\pi ft'}$.

In contrast, the WT of $x(t)$ at a particular time-scale coordinate (t, a) is obtained by a slightly different procedure:

1. Shift the analyzing wavelet function (AWF) $\gamma(t')$ by time t .
2. Scale the time-shifted AWF in time by $\frac{1}{a}$ and normalize the result by $\sqrt{\frac{1}{|a|}}$.
3. Take the inner product of the signal with the time-shifted, time-scaled AWF: $WT_x(t, a) = \langle x, \gamma_{(t,a)} \rangle$, where $\gamma_{(t,a)}(t') = \sqrt{\frac{1}{|a|}} \gamma\left(\frac{t'-t}{a}\right)$.

Like the STFT, the WT suffers from a resolution tradeoff. Because of the difference in the definitions of the two TFRs, the resolution tradeoff manifests itself differently in the WT: at high frequencies, the WT has good time resolution but poor frequency resolution; at low frequencies the WT has good frequency resolution but poor time resolution [HBB92]. Thus, the resolution of the WT does not depend on the choice of analyzing wavelet in the manner that the resolution of the STFT depends on the window function. Instead, the WT's resolution depends on the region of the time-frequency (or time-scale) plane at which the WT is evaluated.

The synthesis of a signal $x(t)$ from its WT is given by

$$x(t) = \frac{c}{a^2} \int_{t'} \int_a WT(t', a) \gamma_{(t',a)}(t) da dt', \quad (2.14)$$

where c is a constant that depends on $\gamma(t')$. The analyzing wavelet must satisfy an *admissibility condition* in order for (2.14) to hold. Daubechies [Dau91] points out that for $\gamma \in L^1(\mathcal{R})$, the admissibility condition is equivalent to

$$\int_{t'} \gamma(t') dt' = 0. \quad (2.15)$$

Thus, $\gamma(t')$ must be bandpass in order to reconstruct the signal from its WT. As noted by Rioul and Vetterli, this implies that $\gamma(t')$ "oscillates in time like a short wave, hence the name 'wavelet' " [RV91].

As with the STFT, the WT is in general complex-valued, so signal analysts often work with the squared magnitude of the WT, known as the **scalogram**:

$$SCAL_x(t, a) = |WT_x(t, a)|^2. \quad (2.16)$$

⁵Here and throughout this work, unless otherwise noted, the inner product is defined as $\langle x, y \rangle = \int x(t)y^*(t) dt$.

The scalogram, which can be considered a quadratic TFR, is similar to the spectrogram in that it provides substantial suppression of crossterms [KBB92]. In this sense, the scalogram is appropriate for use in transient signal detection problems. However, the bandpass nature of the analyzing wavelet required by (2.15) usually implies that the analyzing wavelet will not be well-matched to many models of transient signals, including those examined in the following chapters. Thus, the WT and the scalogram are not considered in the majority of this dissertation, although the concluding chapter addresses the question of extending the methods and results presented in the dissertation to transient signal detection via the WT.

2.2 Overview of the Gabor Transform

2.2.1 Continuous-time Gabor transform

Properties of the continuous-time Gabor transform

Consider again the synthesis equation in (2.10) for the reconstruction of a signal $x(t)$ from its STFT:

$$x(t) = \int_{t'} \int_{f'} STFT_x(t', f') g(t - t') e^{j2\pi f' t} dt df'. \quad (2.17)$$

If the integrals in (2.17) are approximated as summations, with the integrand sampled every Δt in time and every Δf in frequency, one obtains

$$x(t) \approx \sum_m \sum_n STFT_x(m\Delta f, n\Delta t) g(t - n\Delta t) e^{j2\pi(m\Delta f)t} \Delta t \Delta f. \quad (2.18)$$

This is essentially the continuous-time Gabor expansion of $x(t)$.

More precisely, the *Gabor expansion* of a signal $x(t)$ is written as follows:

$$x(t) = \sum_m \sum_n C_{mn} g(t - \alpha n) e^{j2\pi\beta m(t - \alpha n)}. \quad (2.19)$$

The parameters α and β correspond to the respective time and frequency sampling intervals Δt and Δf in (2.18) above. The complex-valued coefficients C_{mn} represent the **Gabor transform** of the signal $x(t)$. Thus, the Gabor transform of a signal is implicitly defined through the decomposition shown in (2.19).

The Gabor expansion can be thought of as a linear combination of the functions

$$g_{mn}(t) = g(t - \alpha n) e^{j2\pi\beta m(t - \alpha n)}, \quad (2.20)$$

which are sometimes termed the *Gabor logons*. Because the Gabor transform is only defined on a lattice in the time-frequency plane, it is important that the Gabor logons be well-localized in the time-frequency plane so that a given coefficient C_{mn} is a good indicator of the signal's time-frequency content at the corresponding location $(n\alpha, m\beta)$ in the time-frequency plane. Particular choices of the window function $g(t)$ that result in well-localized Gabor logons include the one-sided exponential window, to be discussed in Chapters 3–5, and the two-sided Gaussian window, to be discussed in Chapter 6.

It is tempting to call the set of Gabor logons $\mathcal{G} = \{g_{mn}\}$ a *basis*, but such a characterization does not hold true in general, as the logons are not necessarily linearly independent⁶. However, expansion of a signal $x(t)$ by (2.19) is guaranteed only if the set \mathcal{G} constitutes a *frame* [Dau90, Dau91], which is defined as follows:

The collection of functions \mathcal{G} in its underlying Hilbert space \mathcal{H} constitutes a **frame** if there exist constants $A > 0$ and $B < \infty$ such that, $\forall h \in \mathcal{H}$,

$$A\|h\|^2 \leq \sum_{m,n} |\langle h, g_{mn} \rangle|^2 \leq B\|h\|^2.$$

A and B are called the *frame bounds*.

Here, the underlying Hilbert space for \mathcal{G} is $L^2(\mathcal{R})$, together with the inner product $\langle h_1, h_2 \rangle = \int h_1(t)h_2^*(t) dt$.

A necessary condition for \mathcal{G} to constitute a frame is that the parameters α and β have product $\alpha\beta \leq 1$ [Dau91]. The case where $\alpha\beta = 1$ is termed *critical sampling*, while the case where $\alpha\beta < 1$ is termed *oversampling*. Sufficient conditions for \mathcal{G} to constitute a frame are discussed by Daubechies [Dau90, Dau91], and these conditions depend on the window function $g(t)$.

Calculation of the continuous-time Gabor transform

The Gabor transform coefficients cannot be calculated as simply as the STFT is in (2.9) above. If one attempts to discretize time and frequency in the STFT analysis equation in a manner similar to (2.18) above, one obtains

$$STFT_x(m, n) = \int_{t'} x(t')g^*(t' - n\Delta t)e^{-j2\pi(m\Delta f)t'} dt' \quad (2.21)$$

⁶If the logons are *not* linearly independent, then the Gabor transform is *redundant*: a given signal may map to more than one set of Gabor coefficients, and each set of coefficients represents a valid Gabor transform of the signal.

(note that the analysis window has been set equal to the synthesis window). The discrete-valued TFR in (2.21) is called the *sampled STFT* [Orr93b], and it can be shown to be a “blurred” version of the Gabor transform [A⁺90, Orr93b], with the amount of blurring directly related to the orthogonality – or lack thereof – of the Gabor logons.

Most of the existing literature addresses the calculation of the continuous-time Gabor transform coefficients only in the critical sampling case ($\alpha\beta = 1$), so critical sampling will be assumed in the discussion that follows. When $\alpha\beta = 1$, the Gabor expansion in (2.19) may be simplified to

$$x(t) = \sum_m \sum_n C_{mn} g(t - Tn) e^{\frac{j2\pi mt}{T}}, \quad (2.22)$$

where $T = \alpha = \frac{1}{\beta}$. The Gabor transform coefficients in this case are commonly calculated in two ways, the *Zak transform method* (ZTM) and the *biorthogonal function method* (BFM) ⁷.

The ZTM makes use of the Zak transform [Jan88]. For a given signal $x(t)$, the critically sampled Gabor transform coefficients are given by

$$C_{mn} = \int_{t'=0}^T \int_{f'=0}^{\frac{1}{T}} \frac{Z_x(t', f')}{Z_g(t', f')} e^{-j2\pi(\frac{mt'}{T} - nf'T)} dt' df', \quad (2.23)$$

where $Z_x(t', f')$ is the Zak transform of $x(t)$, evaluated in accordance with the critical sampling interval T :

$$Z_x(t', f') = \sum_k x(kT + t') e^{-j2\pi kT f'}. \quad (2.24)$$

The BFM generates a function⁸ $\gamma(t)$, called the *biorthogonal function*, that satisfies the following biorthogonality property:

$$\begin{aligned} \langle g_{mn}, \gamma_{pq} \rangle &= \int_{-\infty}^{\infty} g(t - nT) \gamma^*(t - qT) e^{\frac{j2\pi(m-p)t}{T}} \\ &= \delta(n - q) \delta(m - p). \end{aligned} \quad (2.25)$$

Since the biorthogonal function $\gamma(t)$ is defined in relation to the window function $g(t)$, it is more appropriate to use the notation $\gamma_g(t)$. If a biorthogonal function $\gamma_g(t)$

⁷Orr details a third method for calculating the Gabor transform coefficients that involves deconvolving the sampled STFT. This method will not be addressed here; the reader is referred to [Orr93b] for details.

⁸The notation $\gamma(t)$ is intentionally the same as for the analysis window of the STFT, because the biorthogonal function acts as an analysis window for the Gabor transform.

exists (and this is *not* always the case, even for some well-behaved window functions [Bal92]), then it acts as an analysis window to calculate the Gabor transform coefficients:

$$C_{mn} = \int_{t'} x(t') \gamma_g^*(t' - nT) e^{-\frac{j2\pi mt'}{T}} dt'. \quad (2.26)$$

The biorthogonal function $\gamma_g(t)$ may be generated via the Zak transform of the window function as follows:

$$\gamma_g(t) = \int_0^1 \frac{df'}{Z_g^*(t, f')}. \quad (2.27)$$

Observation of equations (2.23) and (2.27) reveals that both the ZTM and BTM will encounter problems in calculating the Gabor transform coefficients if the Zak transform of the window function contains zeros in the respective regions of integration. It is well known that this can occur even for certain well-behaved window functions [A⁺90, Bal92, Orr93b].

If only a finite subset of the Gabor transform coefficients need to be determined, there is a third method of calculation, which will be termed here the *normal equations method* (NEM). First, recall the Hilbert space \mathcal{H} defined above for the set $\mathcal{G} = \{g_{mn}\}_{m,n=-\infty}^{\infty}$ of Gabor logons, and consider the closed subspace \mathcal{H}_1 of \mathcal{H} spanned by the *finite* subset of Gabor logons $\{g_p\}_{p=1}^P$ (the indices m and n have been reordered into a single index p). Then, for any function $x \in \mathcal{H}$, the *unique* function $\hat{x} \in \mathcal{H}_1$ that best approximates x (in the sense of minimizing $\|x - \hat{x}\|$) must satisfy the property that the *error* function $x - \hat{x}$ is *orthogonal* to \mathcal{H}_1 [BS90, FR94, NS82], i.e.,

$$\langle x - \hat{x}, g_p \rangle = 0, \quad p = 1, \dots, P. \quad (2.28)$$

The latter fact is sometimes called the *orthogonality principle* for Hilbert spaces. Writing \hat{x} as $\sum_{p=1}^P C_p g_p$, one can rewrite (2.28) using the bilinearity of the inner product as

$$\langle x, g_p \rangle = C_1 \langle g_1, g_p \rangle + C_2 \langle g_2, g_p \rangle + \dots + C_P \langle g_P, g_p \rangle, \quad p = 1, \dots, P. \quad (2.29)$$

This gives a set of P equations, called *normal equations*, with P unknowns (the Gabor transform coefficients C_p), and the Gabor transform coefficients can thus be obtained by solving this linear system of equations⁹. The discrete-time equivalent to the NEM, to be described in the next subsection, will play a major role in this dissertation.

⁹Again, if the Gabor logons are not linearly independent, then the linear system of equations will be underdetermined and there will be an infinite number of solutions. However, when only a finite subset of Gabor logons is considered, it is almost always the case that the logons *are* linearly independent; in the latter case, the NEM produces a unique solution.

2.2.2 Discrete-time Gabor transform for finite-length signals

In practice, one must work with finite-length, discrete-time signals. It is thus important to understand how the properties of the Gabor transform and the calculation of the Gabor transform coefficients change in the case of finite-length, discrete-time signals.

Consider first the discretized version of the Gabor expansion of a continuous-time signal $x(t)$, sampled every Δ seconds:

$$x(k\Delta) = \sum_m \sum_n C_{mn} g(k\Delta - n\alpha) e^{j2\pi m\beta(k\Delta - n\alpha)}, \quad k = \dots, -2, -1, 0, 1, 2, \dots \quad (2.30)$$

For finite-duration signals $x(\cdot)$, it can be assumed that $k \in \{0, 1, 2, \dots, L-1\}$, where L is the *length* of the finite signal $x(\cdot)$. As mentioned in the discussion above on the normal equations method (NEM), it is also common to assume that only a finite subset of the Gabor transform coefficients are nonzero. In this case, the expansion (2.30) can be rewritten as

$$x(k\Delta) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g(k\Delta - n\alpha) e^{j2\pi m\beta(k\Delta - n\alpha)}, \quad k = 0, 1, 2, \dots, L-1, \quad (2.31)$$

where M and N are the maximum allowable frequency and time indices, respectively. Traditionally, the discrete-time signal $x(k\Delta)$ is denoted as $x[k]$, with the sampling interval Δ suppressed.

Starting from (2.31), there are three ways to proceed in order to calculate the discrete-time Gabor transform coefficients C_{mn} of $x[k]$:

1. Treat (2.31) as an approximation of the continuous-time Gabor expansion (2.19) and calculate the discrete-time Gabor transform coefficients using a discrete-time biorthogonal function $\gamma_g[k]$ that is the sampled version of the corresponding continuous-time biorthogonal function $\gamma_g(t)$ (if the latter exists). This method was first explored in [FP92] and [PF92].
2. Rewrite (2.31) in a matrix equation of the form $\vec{x} = \mathbf{G}\vec{c}$, and calculate the discrete-time Gabor transform coefficients by finding the vector \hat{c} that minimizes the squared error $\|\mathbf{G}\hat{c} - \vec{x}\|$. This is the discrete-time equivalent to the normal equations method detailed above. The underlying Hilbert space is the vector space \mathcal{R}^L with inner product $\langle \vec{a}, \vec{b} \rangle = \sum_{k=0}^L a_k b_k^*$. This method was first proposed for use with the discrete-time Gabor transform in [BS90] and [EK91].

3. Create a quasi-periodic signal $\tilde{x}[k] = \sum_l x[k + lL]$ from the finite-length signal $x[k]$ and derive a “discrete Gabor transform” (DGT) similar to the way the DFT for finite-duration signals is derived from the Fourier series. The biorthogonality condition (2.25) then has a discrete-time equivalent that can be solved numerically. Once the discrete-time biorthogonal function is found, the DGT coefficients can be calculated by means of inner products in a manner similar to (2.26). The Gabor transform coefficients in this case are doubly periodic, with period M in the frequency direction and period N in the time direction. This method has been explored extensively in the recent literature [Orr93b, QC93, WR90].

Method (1) above, which will be termed here the *discretized biorthogonal function method* (DBFM), will be examined in detail in Chapter 3; method (2), called here the *least squares method* (LSM), will be used throughout Chapters 3 through 5; and method (3), called here the *periodized biorthogonal function method* (PBFM), will be treated in detail in Chapter 4.

For the finite Gabor transform of finite-length, discrete-time signals, the question of the completeness of the set \mathcal{G} of Gabor logons is simple to answer. All three of the methods mentioned above can be rewritten in matrix form as

$$\vec{x} = \mathbf{G}\vec{c}, \quad (2.32)$$

where \vec{x} is $L \times 1$, \mathbf{G} is $L \times MN$, and \vec{c} is $MN \times 1$. Method (3) above assumes that $MN \geq L$, i.e., that the number of transform coefficients is greater than or equal to the number of signal samples. In this case, the system (2.32) is underdetermined and the completeness of \mathcal{G} is assured unless \mathbf{G} is less than full rank. Methods (1) and (2) make no assumptions on the relationship of MN to L , but it is often assumed for transient signal detection applications that $MN \leq L$. In this case, the system (2.32) is overdetermined and \vec{x} *cannot* be reconstructed exactly unless it resides in the MN -dimensional subspace spanned by the Gabor logons.

The Gabor Transform In Transient Signal Detection

3.1 Introduction

This chapter applies the discrete-time Gabor transform to the problem of detecting finite-duration, discrete-time signals. The first part of the chapter, comprising sections 3.2–3.6, examines transient signal detection when the signal components to be detected are assumed to have *known locations* in the time-frequency plane. In section 3.2, a framework for the detection of transient signals, originally developed by Porat and Friedlander [FP92, PF92], is described. Section 3.3 then examines the calculation of the discrete-time Gabor transform coefficients within the Porat-Friedlander framework using two of the methods described at the end of Chapter 2, the discretized biorthogonal function method (DBFM) and the least-squares method (LSM). In this section, *critical sampling* is assumed.

Section 3.4 presents a hypothesis test for the Porat-Friedlander framework and reviews the derivation and analysis of the GLRT detector for this hypothesis test, originally given in [FP89]. Again, the GLRT detector is derived under the assumption that the signal components to be detected have *known* locations in the time-frequency plane. The original contributions of this chapter begin in section 3.5, which examines the calculation of the discrete-time Gabor transform coefficients in the *oversampled* case using the LSM. The corresponding GLRT detector based on the oversampled LSM-Gabor transform is derived. In section 3.6, this last detector based on the oversampled Gabor transform is compared to four detectors based on critically sampled linear transforms:

- the discrete-time Gabor transform calculated via the DBFM

- the discrete-time Gabor transform calculated via the LSM
- the sampled short-time Fourier transform (SSTFT) with analysis window equal to the synthesis window
- the SSTFT with rectangular analysis window.

The comparison is made by examining the robustness of each detector in the presence of various types of signal mismatch.

The second part of the chapter, comprising sections 3.7–3.10, examines transient signal detection using the oversampled LSM-Gabor transform when the signal components to be detected are assumed to have *unknown locations* in the time-frequency plane. Section 3.7 derives the GLRT detector within the Porat-Friedlander framework in the “unknown location” (UL) case and examines the probability distribution of this unknown location detector. Section 3.8 calculates the performance curves of the UL detector and compares the performance of the UL detector to that of the corresponding “known location” (KL) detector of section 3.5. The comparison is made by examining the robustness of each detector in the presence of mismatch in signal component location. Section 3.9 discusses a generalization of the results of the UL case recently studied by Scharf and Friedlander, called *matched subspace detection* [SF94], and it explains how the results of sections 3.7–3.8 are a special case of their analysis.

Finally, section 3.10 summarizes the conclusions of this chapter and presents some ideas for future work.

3.2 A Framework for Transient Signal Detection

This section describes a framework recently developed by Porat and Friedlander for transient signal detection. For brevity, the Porat-Friedlander framework will be referred to henceforth as the PF framework. The PF framework assumes that the finite-length, discrete-time transient signal to be detected can be represented as a finite linear combination of Gabor logons g_{mn} , critically sampled with parameters $\alpha = \beta = 1$:

$$x[k] = x(k\Delta) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g(k\Delta - n) e^{j2\pi mk\Delta} \quad (3.1)$$

$$= \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g_{mn}(k\Delta), \quad k = 0, 1, \dots, L-1, \quad (3.2)$$

where Δ is the sampling interval, L is the signal length, and

$$g_{mn}[k] = g_{mn}(k\Delta) = g(k\Delta - n)e^{j2\pi mk\Delta}. \quad (3.3)$$

The complex-valued constants C_{mn} in (3.2) represent the critically sampled Gabor transform coefficients of $x[\cdot]$.

The window function $g[\cdot]$, from which the Gabor logons $g_{mn}[\cdot]$ in (3.3) above are derived, is assumed to be the sampled version of a one-sided exponential window. The continuous-time version $g(t)$ of the one-sided exponential window is given as follows:

$$g(t) = \sqrt{2\lambda}e^{-\lambda t}u(t), \quad (3.4)$$

where $\lambda > 0$ is the decay parameter of the window. The factor $\sqrt{2\lambda}$ is chosen so that $\int_{-\infty}^{\infty} g(t) dt = 1$. The real part of a Gabor logon with one-sided exponential window, evaluated at $m = 2, n = 2$, and $\lambda = 1$, is shown in Figure 3.1:




Figure 3.1: The real part of a time-shifted and frequency-modulated one-sided exponential window, sampled every $\Delta = \frac{1}{32}$ s. Parameter values are $m = 2, n = 2, \lambda = 1$.

Friedlander and Porat chose the one-sided exponential window as their model for transient signals because it “represents quite well the jump discontinuity and

the gradual decay ... typical to many physical transient phenomena" [FP89]. The discontinuity in the one-sided exponential window indicates that the time index 'n' of a particular Gabor logon $g_{mn}[\cdot]$ corresponds to the *arrival time* of that logon.

The PF framework assumes, then, that the transient signal is a linear combination of a finite number of Gabor logons, where each of the logons is a time-shifted and/or frequency-modulated version of the one-sided exponential window shown in Figure 3.1. Since the number of Gabor logons is finite, the double summation in (3.2) can be reordered into a single summation by means of an appropriate transformation such as

$$l = nM + m \quad (3.5)$$

(so that $n = \lfloor \frac{l}{M} \rfloor$ and $m = l \bmod M$). After the reordering, (3.2) can be rewritten as

$$x[k] = \sum_{l=0}^{MN-1} C_l g_l[k], \quad (3.6)$$

where the sampling interval Δ has been suppressed and where $C_l \mapsto C_{mn}$ and $g_l[k] \mapsto g_{mn}[k]$ in the manner described above. The discrete-time relation in (3.6) can be rewritten in matrix form as

$$\vec{x} = \mathbf{G}\vec{c}, \quad (3.7)$$

where the l th column of \mathbf{G} contains the samples of g_l . In (3.7), \vec{x} is $L \times 1$, \mathbf{G} is $L \times MN$, and \vec{c} is $MN \times 1$.

The PF framework then assumes that the signal \vec{x} is embedded in additive white Gaussian noise and is observed as a vector \vec{y} :

$$\vec{y} = \vec{x} + \vec{v} = \mathbf{G}\vec{c} + \vec{v}. \quad (3.8)$$

Porat and Friedlander originally assumed the signal \vec{x} to be real-valued, so they also assumed real-valued noise with $\vec{v} \sim \mathcal{N}(\vec{0}, \mathbf{I})$. Here, the signal \vec{x} is assumed to be complex-valued in general, so the noise vector \vec{v} is assumed to be $\vec{v} = \vec{v}_{re} + j\vec{v}_{im}$, with $\vec{v}_{re} \sim \mathcal{N}(\vec{0}, \mathbf{I})$, $\vec{v}_{im} \sim \mathcal{N}(\vec{0}, \mathbf{I})$, and $\vec{v} \sim \mathcal{N}(\vec{0}, 2\mathbf{I})$. (Clearly, it is also assumed that \vec{v}_{re} and \vec{v}_{im} are independent.)

A further critical assumption in the PF framework is that the number of possible nonzero signal components (and, therefore, the number of possible nonzero Gabor coefficients) MN is *less than* the length L of the signal \vec{x} (which is also the length of each of the Gabor logons \vec{g}_l and the length of the observation \vec{y}). This is assumed because the detector that eventually arises from the PF framework is based not on the observation vector \vec{y} but on the (estimated) transform coefficient vector \hat{c} . It is only useful to transform the observation vector \vec{y} for detection purposes if the transform reduces the dimensionality of the problem, and the latter will only hold

true if the signal \vec{x} itself is characterized by fewer Gabor transform coefficients than time samples. In general, the assumption $MN < L$ (which will be referred to from here on as assumption A1) will be an accurate one when the actual signal components of \vec{x} are well-modeled by the Gabor logons in (3.3).

Because assumption A1 may not always hold true, it is necessary to incorporate a mismatch vector \vec{e} into the PF framework in order to account for those components in \vec{x} belonging to the orthogonal complement of the subspace spanned by the MN Gabor logons. Thus, $\vec{x} = \mathbf{G}\vec{c} + \vec{e}$ and

$$\vec{y} = \mathbf{G}\vec{c} + \vec{e} + \vec{v}, \quad (3.9)$$

with $\vec{e}^H \mathbf{G} = \mathbf{0} \ (1 \times MN)^1$.

3.3 Calculating the Gabor Transform Within the PF Framework

Within the PF framework, two methods mentioned at the end of Chapter 2 are effective in calculating the discrete-time Gabor transform coefficients: the discretized biorthogonal function method (DBFM) and the least-squares method (LSM).

3.3.1 Discretized biorthogonal function method

As mentioned in Chapter 2, the DBFM uses a discretized version of the continuous-time biorthogonal function $\gamma_g(t)$ to calculate the discrete-time Gabor transform coefficients by means of inner products. In [FP89], Friedlander and Porat showed that for $g(t)$ equal to the one-sided exponential window given in (3.4), the corresponding biorthogonal function $\gamma_g(t)$ (generated by the Zak transform $Z_g(t, f')$ as shown in (2.27)) is given by

$$\gamma_g(t) = \frac{e^{\lambda t}}{\sqrt{2\lambda}} [-u(t+1) + 2u(t) - u(t-1)] . \quad (3.10)$$

Sampling (3.10) every Δ seconds gives

$$\gamma_g[k] = \gamma_g(k\Delta) = \frac{e^{\lambda k\Delta}}{\sqrt{2\lambda}} \{-u(k\Delta+1) + 2u(k\Delta) - u(k\Delta-1)\} . \quad (3.11)$$

¹ \vec{e}^H stands for the complex-conjugate transpose of \vec{e} .

Given the discretized biorthogonal function (3.11), the Gabor transform coefficients are generated by calculating inner products between the observed signal $y[k]$ and time-shifted, frequency-modulated versions of (3.11):

$$C_{mn} = \sum_{k=0}^{L-1} y[k] \gamma_{g,mn}^*[k], \quad (3.12)$$

where

$$\gamma_{g,mn}[k] = \gamma_{g,mn}(k\Delta) = \gamma_g(k\Delta - n) e^{j2\pi mk\Delta}. \quad (3.13)$$

Note that (3.13) assumes the fundamental Gabor expansion parameters α and β are both 1, in agreement with original assumptions of the PF framework. After a reordering as in (3.5) above, (3.12) can be rewritten as the matrix equation

$$\hat{c} = \Gamma \vec{y} = \Gamma G \vec{c} + \Gamma(\vec{e} + \vec{v}), \quad (3.14)$$

where the l th column of Γ contains $\gamma_{g,l}[\cdot]$, sampled in the same way that $x[\cdot]$ is.

The matrix Γ must have a number of columns equal to L (the signal length), but there is no restriction on the number of rows in Γ . However, since it is assumed that the maximum allowable time and frequency shifts for the components of $x[\cdot]$ are M and N , respectively, it is reasonable to use this knowledge to make Γ $MN \times L$. In this case, the biorthogonality condition is rewritten in matrix form as

$$\Gamma G = I (MN \times MN). \quad (3.15)$$

If (3.15) is true, then Γ is said to be a *left-inverse* of G . Since G is non-square with a greater number of rows than columns, it has an infinite number of left-inverses, corresponding to different methods of estimating the Gabor transform coefficients.

3.3.2 Least squares method

As also mentioned in Chapter 2, the LSM begins with the matrix equation (3.9) and seeks to find the coefficient vector \hat{c} that minimizes the squared error $\|G\hat{c} - \vec{y}\|$. It is well known [Sch91] that this “least-squared-error” solution is

$$\hat{c} = (G^H G)^{-1} G^H \vec{y}. \quad (3.16)$$

The matrix $B = (G^H G)^{-1} G^H$, known as the *pseudoinverse* of G , is clearly another left inverse of G .

3.3.3 General transforms within the PF framework

In general, a linear time-frequency transformation of the observation vector \vec{y} can be represented as a premultiplication by an appropriately-dimensioned matrix \mathbf{R} , whose columns contain the “basis” vectors of the transform. Thus, the generic transform step in the PF framework can be written as

$$\begin{aligned}\vec{z} &= \mathbf{R}\vec{y} \\ &= \mathbf{R}\vec{x} + \mathbf{R}\vec{e} + \mathbf{R}\vec{v} \\ &= \mathbf{R}\mathbf{G}\vec{c} + \mathbf{R}\vec{e} + \vec{w},\end{aligned}\tag{3.17}$$

where \vec{z} is the transformed observation vector and $\vec{w} \sim \mathcal{N}(\vec{0}, 2\mathbf{R}\mathbf{R}^H)$.

Ideally, \mathbf{R} will be a left-inverse of \mathbf{G} . Any method of calculating the Gabor transform will correspond to a particular left-inverse of \mathbf{G} (for example, the matrices $\mathbf{\Gamma}$ and \mathbf{B} above). However, the PF framework allows the use of transforms whose corresponding \mathbf{R} matrices are *not* left-inverses of \mathbf{G} . The analysis in subsequent sections of this chapter will make use of two versions of the sampled short-time Fourier transform (SSTFT) whose corresponding \mathbf{R} matrices are indeed *not* left-inverses of \mathbf{G} . The first of these uses an analysis window based on a rectangular window with 1-second duration,

$$g(t) = u(t) - u(t - 1),\tag{3.18}$$

so that the corresponding transform matrix \mathbf{S}_r contains time-shifted and frequency modulated versions of the discrete-time analog of (3.18). The second version of the SSTFT uses an analysis window equal to the synthesis window, i.e., the one-sided exponential window with parameter λ . In this case, the corresponding transform matrix $\mathbf{S}_e = \mathbf{G}^H$. Again, neither \mathbf{S}_r nor \mathbf{S}_e are left-inverses of \mathbf{G} .

3.4 Transient Detection Within the PF Framework

3.4.1 Hypothesis test and GLRT statistic

Consider once again the representation of the transformed observation \vec{z} within the PF framework, assuming a generic transform matrix \mathbf{R} :

$$\vec{z} = \mathbf{R}\mathbf{G}\vec{c} + \mathbf{R}\vec{e} + \vec{w},\tag{3.19}$$

where \mathbf{R} is $MN \times L$ and \vec{z} and \vec{w} are both $MN \times 1$. In addition to assumption A1 above that the number of possible nonzero signal components is small compared to the number of signal samples ($MN < L$), the PF framework makes a further assumption – to be called A2 – that the *actual* number of nonzero signal components is small compared to the number of *possible* nonzero signal components. Assumption A2 is also called the *sparseness* assumption [PF92], and it is made because transient signals are defined to comprise only a small portion of the observation interval. In other words, if relatively many of the components of \vec{c} in (3.19) are nonzero, then the signal \vec{x} would no longer be a transient signal.

To implement assumption A2, Porat and Friedlander postulated that the signal \vec{x} to be detected has a coefficient vector \vec{c} with K nonzero components ($K \ll MN$), all of which have *known locations* within \vec{c} (i.e., known locations in the time-frequency plane). The known-location assumption will be called A3. Since the nonzero component locations are assumed to be known, \vec{c} can always be rearranged so that the K nonzero components reside at the *top* of the coefficient vector. Thus, it is assumed without loss of generality that \vec{c} is of the following form:

$$\vec{c} = \begin{bmatrix} \vec{c}_1 \\ \vec{0} \end{bmatrix}, \quad (3.20)$$

where \vec{c}_1 is $K \times 1$. Of course, all the other components in the framework (3.19) must also be rearranged to conform to the format of \vec{c} in (3.20). For example, \mathbf{G} should be arranged into $[\mathbf{G}_1 \ \mathbf{G}_2]$, where \mathbf{G}_1 is comprised of the K Gabor logons whose time-frequency locations correspond to the (known) locations of \vec{c}_1 . The reordered version of (3.19), then, is given by

$$\begin{aligned} \vec{z} &= \begin{bmatrix} \vec{z}_1 \\ \vec{z}_2 \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{R}_1 \\ \mathbf{R}_2 \end{bmatrix} [\mathbf{G}_1 \ \mathbf{G}_2] \begin{bmatrix} \vec{c}_1 \\ \vec{0} \end{bmatrix} + \begin{bmatrix} \mathbf{R}_1 \vec{e} \\ \mathbf{R}_2 \vec{e} \end{bmatrix} + \vec{w} \\ &= \begin{bmatrix} \mathbf{R}_1 \mathbf{G}_1 \vec{c}_1 \\ \mathbf{R}_2 \mathbf{G}_1 \vec{c}_1 \end{bmatrix} + \begin{bmatrix} \mathbf{R}_1 \vec{e} \\ \mathbf{R}_2 \vec{e} \end{bmatrix} + \vec{w}. \end{aligned} \quad (3.21)$$

When \mathbf{R} is a left-inverse of \mathbf{G} , $\mathbf{R}_1 \mathbf{G}_1 = \mathbf{I}$ and $\mathbf{R}_2 \mathbf{G}_1 = \mathbf{0}$. The dimensions of the matrices in (3.21) are as follows: \mathbf{G}_1 and \mathbf{R}_1^H are $L \times K$, while \mathbf{G}_2 and \mathbf{R}_2^H are $L \times (MN - K)$.

Given (3.20) with assumptions A1, A2, and A3, detection of the transient signal \vec{x} within the PF framework is tested by the following binary hypothesis test:

$$\mathbf{H}_0 : \vec{c}_1 = \vec{0}$$

vs.

$$\mathbf{H}_1 : \vec{c}_1 \neq \vec{0}. \quad (3.22)$$

The alternative hypothesis here is a *composite* hypothesis test, so Porat and Friedlander proposed the use of the generalized likelihood ratio test (GLRT) to derive a detection statistic. The GLRT for a simple null hypothesis and a composite alternative hypothesis, given observations \vec{z} and parameters \vec{c}_1 , is defined as

$$\frac{\max_{\vec{c}_1} p_1(\vec{z} | \vec{c}_1)}{p_0(\vec{z})}, \quad (3.23)$$

where $p_i(\vec{z})$ is the probability distribution of \vec{z} under hypothesis i . Taking the natural log of (3.23) gives a more suitable form of the GLRT:

$$t = \max_{\vec{c}_1} \{2 \log p_1(\vec{z} | \vec{c}_1)\} - 2 \log p_0(\vec{z}), \quad (3.24)$$

where the factor of 2 is included for convenience.

Since the detector does not have knowledge of the mismatch vector \vec{e} , the GLRT statistic is calculated assuming $\vec{e} = \vec{0}$ (although \vec{e} reappears when calculating the probability distributions of the GLRT statistic for performance analysis). Using this assumption, and recalling that $\vec{w} \sim \mathcal{N}(\vec{0}, 2\mathbf{R}\mathbf{R}^H)$ in (3.19), it is clear that the observation vector $\vec{z} \sim \mathcal{N}(\mathbf{R}\mathbf{G}\vec{c}_1, 2\mathbf{R}\mathbf{R}^H)$. Friedlander and Porat used the above information to show in [FP89] that, for transforms \mathbf{R} that are left inverses² of \mathbf{G} (i.e., $\mathbf{R}\mathbf{G}=\mathbf{I}$), the GLRT statistic for the hypothesis test of (3.22) is given by

$$t = \vec{z}^H \mathbf{\Lambda}^{-1} \vec{z} - \vec{z}_2^H \mathbf{\Lambda}_{22}^{-1} \vec{z}_2, \quad (3.25)$$

where

$$\mathbf{R}\mathbf{R}^H = \mathbf{\Lambda} = \begin{bmatrix} \mathbf{\Lambda}_{11} & \mathbf{\Lambda}_{12} \\ \mathbf{\Lambda}_{21} & \mathbf{\Lambda}_{22} \end{bmatrix}. \quad (3.26)$$

It was shown in [PF92], which assumed real-valued noise, that the test statistic t is χ^2 distributed with K degrees of freedom, with the distribution being central χ^2 under hypothesis H_0 and noncentral χ^2 under hypothesis H_1 , with noncentrality parameter

$$\nu = (\mathbf{G}_1 \vec{c}_1 + \vec{e})^H \mathbf{P} (\mathbf{G}_1 \vec{c}_1 + \vec{e}). \quad (3.27)$$

²The calculations become intractable when \mathbf{R} is not a left-inverse of \mathbf{G} .

The matrix \mathbf{P} in (3.27) is a projection matrix and is given by

$$\begin{aligned} \mathbf{P} &= (\mathbf{R}_1^H - \mathbf{R}_2^H(\mathbf{R}_2\mathbf{R}_2^H)^{-1}\mathbf{R}_2\mathbf{R}_1^H) \\ &\times (\mathbf{R}_1(\mathbf{I} - \mathbf{R}_2^H(\mathbf{R}_2\mathbf{R}_2^H)^{-1}\mathbf{R}_2)\mathbf{R}_1^H)^{-1} \\ &\times (\mathbf{R}_1 - \mathbf{R}_1\mathbf{R}_2^H(\mathbf{R}_2\mathbf{R}_2^H)^{-1}\mathbf{R}_2). \end{aligned} \quad (3.28)$$

A more compact (but less understandable) form of \mathbf{P} is given in [PF92] using projection operators.

In the case where complex-valued noise is used, the test statistic t is still χ^2 -distributed, with a central distribution under \mathbf{H}_0 and a noncentral distribution with noncentrality parameter ν given by (3.27) under H_1 . However, there are $2K$ degrees of freedom in the χ^2 distribution for either hypothesis, because there are two random noise components in the complex noise per signal sample.

3.4.2 Analysis of GLRT statistic within PF framework

Note that when \vec{c} takes the form in (3.20) and when \mathbf{R} is a left-inverse of \mathbf{G} , as was assumed in deriving the GLRT statistic (3.25), the signal \vec{x} is given by $\vec{x} = \mathbf{G}_1\vec{c}_1 + \vec{e}$. Thus, the noncentrality parameter in (3.27) can be rewritten

$$\nu = \vec{x}^H \mathbf{P} \vec{x}. \quad (3.29)$$

Now for a given false alarm probability and for a given number of degrees of freedom, the detection probability of the test statistic t for the hypothesis test (3.22) is a monotonically increasing function of the noncentrality parameter ν . It is clear from (3.29), however, that ν depends on the signal \vec{x} itself, which in turn depends on the mismatch vector \vec{e} . Thus, the detection performance of t is a function of the energy $\vec{e}^H \vec{e}$ of the mismatch vector \vec{e} in relation to the energy $E = \vec{x}^H \vec{x} = (\mathbf{G}_1\vec{c}_1 + \vec{e})^H (\mathbf{G}_1\vec{c}_1 + \vec{e})$ of the entire signal \vec{x} . This relationship is quantified by the variable

$$\rho = \frac{\vec{e}^H \vec{e}}{E}, \quad (3.30)$$

which measures the fraction of mismatch energy to total signal energy.

The robustness of t in the presence of mismatch is best measured by the *minimum* value of the noncentrality parameter in (3.27), denoted as ν_{\min} , minimized over all vectors \vec{e} satisfying the previously established constraints

$$\begin{aligned} \vec{e}^H \mathbf{G}_1 \vec{c}_1 &= 0 && (\text{orthogonality}) \\ \vec{e}^H \vec{e} &= E - \vec{c}_1^H \mathbf{G}_1^H \mathbf{G}_1 \vec{c}_1 && (\text{energy}). \end{aligned} \quad (3.31)$$

The minimum noncentrality parameter represents the worst-case GLRT detector performance, so a relatively high value of ν_{min} indicates a relatively robust GLRT detector. Note that the orthogonality constraint in (3.31) is slightly weaker than the original condition that $\vec{e} \in \mathbf{G}_1^\perp$: here, $\vec{e} \in \mathbf{G}_1^\perp$.

It was shown in [PF92] that the solution to the minimization problem

$$\min_{\vec{e}} \nu = \min_{\vec{e}} (\mathbf{G}_1 \vec{c}_1 + \vec{e})^H \mathbf{P} (\mathbf{G}_1 \vec{c}_1 + \vec{e}), \quad (3.32)$$

subject to the constraints on \vec{e} in (3.31), is given by

$$\nu_{min} = E \left(\sqrt{(1-\rho)\phi} - \sqrt{(1-\phi)\rho} \right)^2, \quad (3.33)$$

where

$$\phi = \frac{\vec{c}_1^H \mathbf{G}_1^H \mathbf{P} \mathbf{G}_1 \vec{c}_1}{\vec{c}_1^H \mathbf{G}_1^H \mathbf{G}_1 \vec{c}_1} \quad (3.34)$$

and where $\rho \leq \phi$. The necessity of the latter inequality was not checked in [PF92] but is verified in Appendix 3.A at the end of this chapter.

Because ν_{min} in (3.33) is dependent on the parameter ϕ , and ϕ in turn is dependent on $\mathbf{G}_1 \vec{c}_1$, the expression for ν_{min} in (3.33) is dependent on the values of the nonzero signal components found in \vec{c}_1 . However, note that when ρ is fixed, ν_{min} is a monotonically increasing function of ϕ when $\rho \leq \phi$. As described by Porat and Friedlander in [PF92], it is actually possible to derive a *signal-independent* lower bound on ϕ . To do this, first observe that ϕ in (3.34) is of the form

$$\phi = \frac{\vec{x}_0^H \mathbf{P} \vec{x}_0}{\vec{x}_0^H \vec{x}_0}, \quad (3.35)$$

where $\vec{x}_0 = \mathbf{G}_1 \vec{c}_1$. Because \mathbf{P} is a projection of the $L \times K$ matrix

$$(\mathbf{R}_1^H - \mathbf{R}_2^H (\mathbf{R}_2 \mathbf{R}_2^H)^{-1} \mathbf{R}_2 \mathbf{R}_1^H)$$

(recall that K is the number of nonzero components in \vec{x} and L is the number of samples in \vec{x}), \mathbf{P} can be decomposed into $\mathbf{P} = \mathbf{U} \mathbf{U}^H$, where \mathbf{U} is $L \times K$ and $\mathbf{U}^H \mathbf{U} = \mathbf{I}$ ($K \times K$). Also, as \vec{x}_0 lies in the span of \mathbf{G}_1 , it can be expressed as $\vec{x}_0 = \mathbf{V} \vec{c}_2$, where the columns of \mathbf{V} ($L \times K$) form an orthonormal basis for \mathbf{G}_1 (so that \mathbf{V} is unitary). Using the last two relations, ϕ can be rewritten as

$$\phi = \frac{\vec{c}_2^H \mathbf{V}^H \mathbf{U} \mathbf{U}^H \mathbf{V} \vec{c}_2}{\vec{c}_2^H \vec{c}_2}. \quad (3.36)$$

In this case, the minimum value of ϕ over all vectors \vec{c}_2 is given by the square of the minimum eigenvalue of $\mathbf{U}^H \mathbf{V}$ [GL83]:

$$\phi_{Min} = \min\{\xi^2 : \mathbf{P}\vec{v} = \xi\vec{v} \text{ for some } \vec{v}\}. \quad (3.37)$$

This ϕ_{Min} (note the distinction in notation between ϕ_{min} and ϕ_{Min}) can then be substituted into (3.33) along with the fixed value of ρ to obtain the signal-independent minimum ν_{Min} of the noncentrality parameter.

The value of ϕ_{Min} in (3.37) leads to the minimum noncentrality parameter ν_{min} over all possible mismatch vectors \vec{e} satisfying (3.31) and over all possible signals \vec{x}_0 corresponding to a given set of K Gabor logons that comprise \mathbf{G}_1 . In order to find the most general minimum of ν , it is necessary to minimize ϕ not only over all \vec{x}_0 corresponding to a *given* set of K component locations, but over all \vec{x}_0 corresponding to *any* set of K component locations. There is no easy way to do this type of minimization; the exact answer requires that one calculate ϕ_{Min} in (3.37) for each of $\binom{MN}{K}$ possible combinations of the K nonzero component locations and then take the minimum among all the calculated ϕ_{Min} values. In [FP92] and [PF92], Porat and Friedlander used simulated annealing to perform this last minimization, but a different method will be used shortly in Section 3.6. Once the “location” minimization is performed, the “global” minimum ϕ_{MIN} can again be substituted along with the fixed value of ρ into (3.33) to obtain what is finally a signal-independent, location-independent lower bound on the noncentrality parameter, denoted as ν_{MIN} .

3.5 PF Analysis Results and Suggestion For Robustness Improvement: Oversampling

3.5.1 Porat-Friedlander results

In [PF92], Porat and Friedlander used their framework to compare the performances of GLRT statistics based on three critically sampled linear TFRs:

- the discrete-time Gabor transform with one-sided exponential window, calculated via the DBFM using the biorthogonal function (3.11). This corresponds to $\mathbf{R} = \mathbf{\Gamma}$ from (3.14), and the GLRT statistic based on $\mathbf{\Gamma}$ will be called the DBFM-Gabor detector.
- the discrete-time Gabor transform with one-sided exponential window, calculated via the LSM. This corresponds to $\mathbf{R} = \mathbf{B}$, the pseudoinverse of \mathbf{G} . The

GLRT statistic based on \mathbf{B} will be called the LSM-Gabor detector.

- the SSTFT with 1-second rectangular window³. This corresponds to $\mathbf{R} = \mathbf{S}_r$ as detailed in subsection 3.3.3, and the GLRT statistic based on \mathbf{S}_r will be called the RW-SSTFT detector.

The comparison was made by determining the most general minimum noncentrality parameter ν_{MN} (as described at the end of the previous section) for several cases that determine the effects of using incorrect information in the formation of the three transform matrices \mathbf{R} above. Incorrect information was also termed “mismatch” in [PF92], but this type of mismatch should be distinguished from the type of mismatch modeled by the mismatch vector \vec{e} in the previous two sections. The latter type of mismatch arises from incomplete modeling of the signal \vec{x} by the signal subspace matrix \mathbf{G} , while the former type of mismatch arises from incorrect assumptions about a *given* signal \vec{x} that affect the formation of the transform matrices \mathbf{R} . To distinguish the two types of mismatch, the following terminology will be used from this point on: *model mismatch* will refer to the mismatch modeled by \vec{e} within the PF framework, while *information mismatch* will refer to the mismatch that causes incorrect formation of the transform matrices \mathbf{R} .

The specific parameter values used by Porat and Friedlander for their analysis were as follows: $M = N = 8$, $K = 6$, $L = 256$, $\Delta = \frac{1}{32}$ (see equation (3.2) and the discussion just before equation (3.20) to review the meanings of these parameters). As mentioned above, Porat and Friedlander worked with *real*-valued Gabor logons instead of complex-valued logons in forming \mathbf{R} and \mathbf{G} . For every possible signal component location in the time-frequency plane, they replaced the complex-valued Gabor logon (see (3.3)) containing a complex exponential term with *two* real-valued logons, one containing a cosine term and the other containing a sine term. Because of this, the χ^2 distribution of the GLRT statistic t for the case of real-valued signals *also* has $2K$ degrees of freedom under either \mathbf{H}_0 or \mathbf{H}_1 (not K degrees of freedom, as reported in [PF92])⁴. Thus, there is essentially no difference, as far as the performance of the GLRT statistic is concerned, between using real-valued signals and using complex-valued signals: the distributions of the statistic in both cases are identical under both \mathbf{H}_0 and \mathbf{H}_1 . It is important to note that the number of degrees of freedom does not affect the analysis of this section or the next, because all the GLRT statistics have the *same* number of degrees of freedom for a given set

³Critically sampling in the SSTFT is equivalent to taking $\Delta t = \Delta f = 1$ in (2.21).

⁴One other consequence of using two real-valued logons per time-frequency location is that K should be replaced everywhere by $2K$ and MN should be replaced everywhere by $2MN$ in section 3.4.

of parameters in the PF framework; it is the noncentrality parameter that varies as the transform \mathbf{R} varies.

There were seven cases of information mismatch examined by Porat and Friedlander in [PF92]:

1. There is no information mismatch (perfect knowledge of the signal).
2. The decay parameter λ in the one-sided exponential window (3.4) is assumed to be 1, but it is actually 0.5.
3. The parameter λ is assumed to be 1, but it is actually 2.
4. The parameter λ is assumed to be 0.5, but it is actually 1.
5. The parameter λ is assumed to be 2, but it is actually 1.
6. The arrival times are assumed to take on integer values (ranging from 0 to 7), but they actually take on values that are odd multiples of 0.5 (ranging from 0.5 to 7.5).
7. The modulation frequencies are assumed to take on integer values (ranging from 0 to 7), but they actually take on values that are odd multiples of 0.5 (ranging from 0.5 to 7.5).

Cases 2–5 will be termed here *shape mismatch*, while cases 6–7 will be termed *location mismatch*. For the last two cases, Porat and Friedlander included *two* logons in the transform matrices \mathbf{R}_1 for each of the logons in \mathbf{G}_1 corresponding to the nonzero signal components, namely the two logons “surrounding” each signal component. For example, if the logon corresponding to $(t, f) = (1, 2)$ was included in \mathbf{G}_1 , the logons corresponding to $(1, 2)$ and $(2, 2)$ were included in \mathbf{R}_1 .

Table 3.1 show the values for ϕ_{MIN} that Porat and Friedlander calculated in [PF92] for the three detectors mentioned above in the seven cases mentioned above. As noted in the previous section, they used simulated annealing to perform the location minimization. The ϕ_{MIN} values in Table 3.1 can be substituted along with various values of ρ (provided $\rho \leq \phi$) into (3.33) to obtain the behavior of ν_{MIN} versus ρ . Recall, though, that ν_{MIN} is a monotonically increasing function of ϕ_{MIN} for fixed ρ such that $\rho \leq \phi$, so the robustness of the three detectors in the presence of information mismatch can be determined directly from Table 3.1: higher values of ϕ_{MIN} (up to a maximum value of 1) mean greater robustness. An examination of Table 3.1, then, reveals the following three points:

Case	Detector		
	Γ	\mathbf{B}	\mathbf{S}_r
1	0.6688	1	0.7107
2	0.6980	0.8397	0.5717
3	0.3273	0.7937	0.5771
4	0.6840	0.8425	0.7107
5	0.3361	0.7935	0.7107
6	0.0476	0.0451	0.1574
7	0.2589	0.2622	0.4137

Table 3.1: Values of ϕ_{MIN} for different cases of information mismatch, as reported in [PF92].

- The LSM-Gabor detector \mathbf{B} is the most robust (i.e., has the highest value of ϕ_{MIN}) of the three detectors when there is no mismatch or any type of shape mismatch.
- The DBFM-Gabor detector is the least robust of the three in almost every case.
- The SSTFT detector is the most robust of the three when there is any type of location mismatch. However, all three detectors perform significantly worse when there is any location mismatch.

Note that the detector based on \mathbf{S}_r is independent of the λ parameter. Thus, the performance of the RW-SSTFT detector in the presence of shape mismatch is determined by how well the 1-second rectangular window matches the actual one-sided exponential window. Inspection of cases 2–5 reveals that the 1-second rectangular window is most well-matched to a one-sided exponential window with $\lambda = 1$.

3.5.2 Oversampling to improve GLRT robustness in the presence of location mismatch

It is clear from the Table 3.1 and the subsequent observations that the most harmful type of information mismatch within the PF framework is location mismatch, any mismatch between the actual and expected time-frequency locations of the signal's logons. One simple way of improving the robustness of any Gabor-transform-based

detector in the presence of location mismatch is to increase the density of the time-frequency lattice over which the Gabor transform is calculated. Recalling the original definition of the discrete-time Gabor expansion,

$$x(k\Delta) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g(k\Delta - n\alpha) e^{j2\pi m\beta(k\Delta - n\alpha)}, \quad k = 0, 1, 2, \dots, L-1, \quad (3.38)$$

this increase in density corresponds to the *oversampling* case where $\alpha\beta < 1$. Throughout the remainder of this dissertation, the oversampled Gabor transform will be implemented with values of $\alpha = \beta = 0.5$, meaning that the density of Gabor logons is increased by a factor of 2 in both the time and frequency directions. This is done for simplicity, and it should quickly be noted that different values of α and β can be used in the following analysis, with minimal adjustments⁵.

The following example illustrates the best-case improvement resulting from oversampling the Gabor transform. First, consider the left-hand side of Figure 3.2, which shows the magnitudes of the critically sampled LSM-Gabor coefficients (assuming one-sided exponential window) for a two-component signal. The two components each have magnitude 1 and are located at (t, f) coordinates of $(1, 1)$ and $(3, 1)$, respectively. There is no mismatch in this case: the arrival times and modulation frequencies of the two components are integer-valued. The corresponding contour plot shows that the two components are well-resolved. Then consider the right-hand side of Figure 3.2, which shows the same thing as the left-hand side except that the (t, f) coordinates of the two signal components are now $(1, 1.5)$ and $(3, 1.5)$ instead. Here, there is a location mismatch in the frequency direction. The corresponding contour plot shows that the two components are much less well-resolved.

Now consider the magnitudes of the oversampled LSM-Gabor transform coefficients of the same two signals, shown in Figure 3.3. To compensate for the fact that the density of the oversampled Gabor logons is twice that of the critically sampled Gabor logons in both the time and frequency directions, the LSM-Gabor coefficients were calculated over a smaller overall range of (t, f) coordinates. In other words, $M = N = 8$ with $\alpha = \beta = 0.5$ in (3.38) means that the maximum arrival time and maximum modulation frequency of any of the oversampled Gabor logons is 3.5 (as compared to 7 for the critically sampled Gabor logons). The oversampled LSM-Gabor transform coefficients are calculated in the same way as the critically sampled LSM-Gabor transform coefficients, using the pseudoinverse. In the oversampled case, of course, the signal subspace matrix \mathbf{G} contains the oversampled Gabor logons instead of the critically sampled Gabor logons (see the next section for the

⁵There are consequences in terms of the stability of the Gabor transform coefficients when α and β are taken to be too small, and this will be discussed in Chapter 4

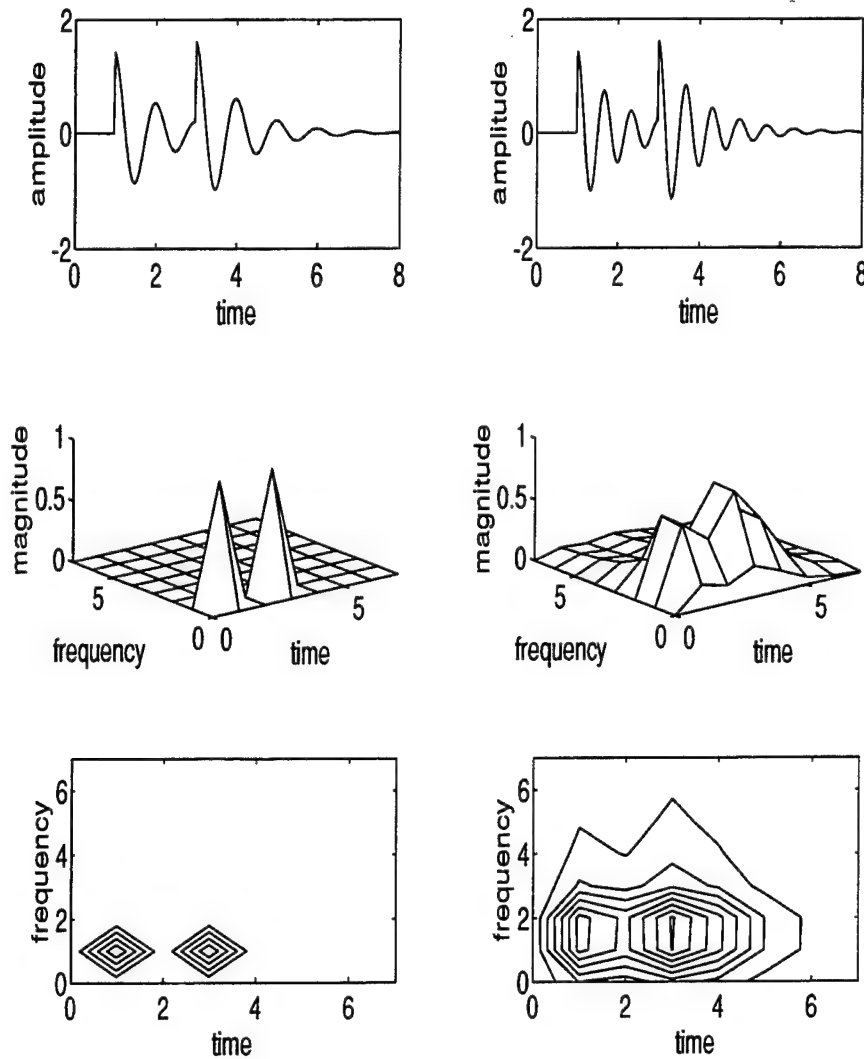


Figure 3.2: Left-hand side: time history (real part), critically sampled LSM-Gabor transform coefficients, and corresponding contour plot for signal with components at (1,1) and (3,1) (**no** location mismatch). Right-hand side: time history (real part), critically sampled LSM-Gabor coefficients, and corresponding contour plot for signal with components at (1,1.5) and (3,1.5) (frequency mismatch of 0.5).

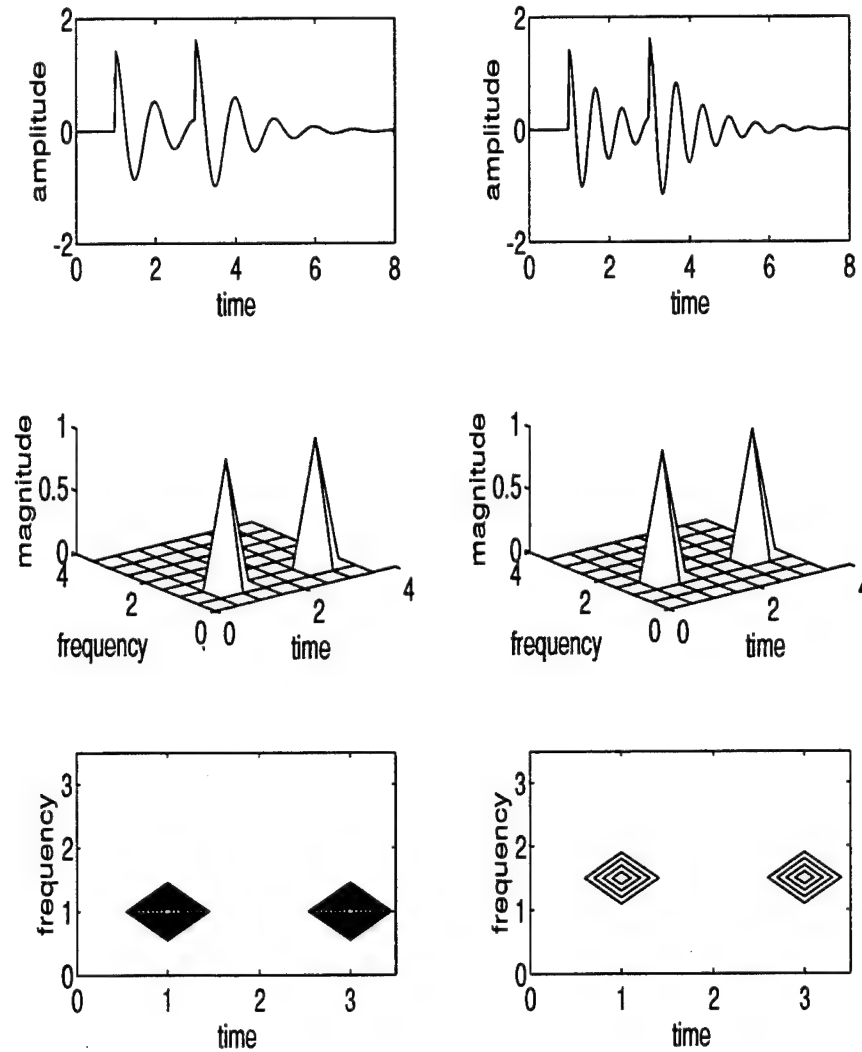


Figure 3.3: Left-hand side: time history (real part), oversampled LSM-Gabor coefficients, and corresponding contour plot for signal with components at (1,1) and (3,1) (no location mismatch). Right-hand side: time history (real part), oversampled LSM-Gabor coefficients, and corresponding contour plot for signal with components at (1,1.5) and (3,1.5) (*also* no location mismatch).

exact meaning of this). Figure 3.3 clearly shows that the two signal components are well-resolved in either case. The reason for this is that *neither* of the signals is information-mismatched with the oversampled Gabor logons. In the best-case, then, the improvement in component resolution due to oversampling is dramatic.

The improvement from oversampling is not always as dramatic as that shown in the above example. In fact, as will be discussed in the next chapter, one encounters stability problems when calculating the oversampled LSM-Gabor transform coefficients. However, the reason that using oversampled Gabor GLRT detector within the PF framework should improve detector robustness in the presence of location mismatch is obvious: with logons that are “closer together” (by a factor of 2), the worst-case location mismatch is only 0.25 in either the time or frequency directions, instead of 0.5. Even in the worst case of mismatch, most of the energy of the signal \vec{x} will be captured by the surrounding transform coefficients. This intuitive line of reasoning is supported by the results of the next section.

3.6 Comparison of Detector Performances: Oversampled Gabor Detector vs. Critically Sampled Detectors

This section compares the robustness of a GLRT detector based on the oversampled LSM-Gabor transform to the robustness of GLRT detectors based on four critically sampled linear transforms. The transform matrices corresponding to the five detectors are as follows:

- $\mathbf{R} = \mathbf{\Gamma}$ representing the discrete-time, critically sampled DBFM-Gabor transform with one-sided exponential window.
- $\mathbf{R} = \mathbf{B}_c$, representing the discrete-time, critically sampled LSM-Gabor transform with one-sided exponential window.
- $\mathbf{R} = \mathbf{S}_r$, representing the critically sampled SSTFT with rectangular window.
- $\mathbf{R} = \mathbf{S}_e = \mathbf{G}^H$, representing the critically sampled SSTFT with one-sided exponential window.
- $\mathbf{R} = \mathbf{B}_o$, representing the discrete-time, oversampled LSM-Gabor transform with one-sided exponential window.

The matrix \mathbf{B}_o for the oversampled LSM-Gabor transform is given by $(\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H$, where the columns of \mathbf{G}_o contain, instead of the critically sampled logons in \mathbf{G} , oversampled logons of the form

$$g_{mn}[k] = g_{mn}(k\Delta) = g(k\Delta - 0.5n)e^{j2\pi(0.5m)(k\Delta - 0.5n)}. \quad (3.39)$$

The original analysis by Porat and Friedlander that was detailed in the previous section used GLRT detectors based on the first three transforms above. The SSTFT with one-sided exponential window is included in the analysis here as an additional basis of comparison.

3.6.1 Modified framework

Because the oversampled LSM-Gabor detector with $\alpha = \beta = 0.5$ looks for signal energy at twice the resolution of any of the three critically sampled detectors in both the time and frequency directions, two assumptions need to be made in order to ensure a fair comparison with the critically sampled detectors.

The first assumption, made to keep the computational load the same for all the detectors, is that M and N in equation (3.38) must be the same for all the detectors. For the oversampled detector, values of $M = N = 8$ signify that the maximum arrival time and maximum modulation frequency for any of the oversampled Gabor logons are both 3.5 (not 7, as with the critically sampled detectors). Thus, the actual time-frequency coordinates (t, f) of the $K = 6$ nonzero signal components are assumed to satisfy $0 \leq t \leq 3.5$ and $0 \leq f \leq 3.5$. The critically sampled transforms are thus calculated over a greater time-frequency range than necessary (0 to 7 in both time and frequency, when the actual components range only from 0 to 3.5), but this does not affect detector performance since the locations of the signal components are assumed known.

The second assumption, made to ensure that the signal to be detected is the same for all the detectors, is that the time-frequency coordinates (t, f) of the $K = 6$ nonzero signal components take on values contained in the critically sampled ($\alpha = \beta = 1$) Gabor transform lattice, i.e., t and f are always integer-valued. Location mismatch is then simulated by shifting from points on the critically sampled lattice. For example, a frequency mismatch of f_m is simulated by adding f_m to the frequency coordinates of all the logons in the signal subspace matrix \mathbf{G} . Figure 3.4 shows the possible signal component locations for the cases of no location mismatch and 0.25 frequency mismatch. Actually, this is exactly how location mismatch was simulated by Porat and Friedlander in the analysis of the previous section, but it needs to be made clear here that location mismatch for the oversampled detector is still simulated




Figure 3.4: Possible signal component locations for modified PF framework, in the cases of no location mismatch (marked by ‘o’) and frequency mismatch of 0.25 (marked by ‘x’).

by shifting logons of the *critically* sampled signal subspace matrix \mathbf{G} , not by shifting the logons of the oversampled signal subspace matrix \mathbf{G}_o .

In the analysis that follows, the values of the PF framework parameters are set to the values that Porat and Friedlander used: $M = N = 8$, $K = 6$, $L = 256$, and $\Delta = \frac{1}{32}$. There are two differences, however, in the methods used to obtain ϕ_{MIN} (and thus, ν_{MIN}). First, it is *not* assumed here that the signals and transform coefficients are necessarily real-valued. Thus, only one complex-valued logon (instead of two real-valued logons) is included in \mathbf{G} and \mathbf{R} for every possible signal component location in the time-frequency plane. Second, simulated annealing is not used to find the location minimum ϕ_{MIN} over all $\binom{16}{6} = 8008$ possible location combinations of the K signals components⁶ Instead, the location minimization is accomplished here in three steps:

1. Generate randomly several hundred location combinations.
2. Observe what “types” of location combinations are likely to produce low values of ϕ_{Min} (the signal-independent minimum for a given location combination).

⁶The total number of possible time-frequency coordinates for the signal components is 16 instead of 64 because of the assumption above regarding computational loads; see Figure 3.4.

3. Examine slight permutations of the "bad" location combinations of step 2 to determine finally a location-independent minimum, ϕ_{MIN} .

This method actually turns out to be an "intelligent" form of simulated annealing. In real-life detection problems, the number of possible location combinations would probably be much greater than 8008, so an exhaustive search for ϕ_{MIN} would likely be impractical; for practical purposes, either simulated annealing or the method described above should be used to find ϕ_{MIN} .

3.6.2 Location mismatch

Table 3.2 below lists the values of ϕ_{MIN} for the five detectors in six cases of location mismatch. Cases 1–3 involve arrival-time mismatches of 0, 0.25, and 0.5, respectively. In these three cases, because an arrival-time mismatch is expected, two logons are included in each transform submatrix \mathbf{R}_1 for every logon in the signal submatrix \mathbf{G}_1 , in such a way that if a logon with (t, f) coordinates (m_0, n_0) is included in \mathbf{G}_1 , logons with (t, f) coordinates (m_0, n_0) and $(m_0, n_0 + 1)$ are included in \mathbf{R}_1 (critically sampled). For the oversampled LSM-Gabor detector, the two logons have (t, f) coordinates (m_0, n_0) and $(m_0, n_0 + 0.5)$ instead. Cases 4–6 involve frequency mismatches of 0, 0.25, and 0.5, respectively. In these three cases, because a frequency mismatch is expected, two logons are again included in each \mathbf{R}_1 for every logon in the signal submatrix \mathbf{G}_1 , this time in such a way that if a logon with (t, f) coordinates (m_0, n_0) is included in \mathbf{G}_1 , logons with (t, f) coordinates (m_0, n_0) and $(m_0 + 1, n_0)$ are included in \mathbf{R}_1 (critically sampled). For the oversampled LSM-Gabor detector, the two logons have (t, f) coordinates (m_0, n_0) and $(m_0 + 0.5, n_0)$ instead. The above procedure is followed for the sake of consistency even when the location mismatch (arrival-time or frequency) is 0. Note that the maximum value for ϕ is 1, which corresponds to matched filter (ideal) detection.

Examination of Cases 1–3 reveals that the oversampled LSM-Gabor detector is clearly the most robust of all the detectors in the presence of arrival-time mismatch. Of particular interest is Case 3, representing an arrival-time mismatch of 0.5 from the critically sampled Gabor lattice depicted in Figure 3.4. This is the worst case of arrival-time mismatch for the critically sampled detectors, but any signal component with an arrival-time mismatch of 0.5 from the critically sampled Gabor lattice is *matched* to the neighboring oversampled Gabor logon. Thus, for the oversampled LSM-Gabor detector, the worst-case arrival-time mismatch is actually 0.25. Even in this case (case 2 in Table 3.2), the detector \mathbf{B}_o is still most robust.

Examination of Cases 4–6 reveals that the oversampled LSM-Gabor detector is

Case	Detector				
	Γ	B_c	S_r	S_c	B_o
1	0.7165	1	0.8615	0.5616	1
2	0.1633	0.1512	0.2285	0.2042	0.4549
3	0.0955	0.0929	0.1555	0.1410	1
4	0.7165	1	0.7348	0.4794	1
5	0.4360	0.5330	0.5636	0.4518	0.8180
6	0.1471	0.1871	0.2496	0.1987	1

Table 3.2: Values of ϕ_{MIN} for different cases of location mismatch.

again clearly the most robust of all the detectors in the presence of frequency mismatch. In general, frequency mismatch is not as harmful as arrival-time mismatch, probably because the discontinuity in the one-sided exponential window makes each of the transforms particularly sensitive to arrival-time mismatch. Figures 3.5–3.8 show the variation of ν_{MIN} versus the fraction of mismatch energy ρ for cases 2, 3, 5, and 6 in Table 3.2 above. The values of ν_{MIN} are calculated from equation (3.33) using the values of ϕ_{MIN} shown in Table 3.2.




Figure 3.5: Detector performance for arrival-time mismatch of 0.25. 'CS' stands for "critically sampled"; 'OS' stands for "oversampled."




Figure 3.6: Detector performance for arrival-time mismatch of 0.5. 'CS' stands for "critically sampled"; 'OS' stands for "oversampled." Note that the curve for the LSM-Gabor (OS) detector resides at the very top of the graph.




Figure 3.7: Detector performance for frequency mismatch of 0.25. ‘CS’ stands for “critically sampled”; ‘OS’ stands for “oversampled.”




Figure 3.8: Detector performance for frequency mismatch of 0.5. ‘CS’ stands for “critically sampled”; ‘OS’ stands for “oversampled.” Note that the curve for the LSM-Gabor (OS) detector resides at the very top of the graph.

3.6.3 Shape mismatch

Table 3.3 below lists the values of ϕ_{MIN} for the five detectors in seven cases of shape mismatch. Case 1 is the no mismatch case: both the actual and assumed values for λ in the one-sided exponential window are 1. Cases 2–4 represent situations where λ is overestimated: in case 2, λ is assumed to be 1 when it is actually 0.5; in case 3, λ is assumed to be 2 when it is actually 1; in case 4, λ is assumed to be 4 when it is actually 1. Cases 5–7 represent situations where λ is underestimated: in case 5, λ is assumed to be 0.5 when it is actually 1; in case 6, λ is assumed to be 1 when it is actually 2; in case 7, λ is assumed to be 1 when it is actually 4. In all of these cases, only one logon is included in \mathbf{R}_1 for every logon in \mathbf{G}_1 .

Case	Detector				
	Γ	\mathbf{B}_c	\mathbf{S}_r	\mathbf{S}_c	\mathbf{B}_o
1	0.7165	1	0.7287	0.3890	1
2	0.7443	0.8515	0.5761	0.2800	0.8767
3	0.4488	0.8123	0.7287	0.3071	0.8222
4	0.1440	0.4194	0.7287	0.1424	0.4191
5	0.7453	0.8487	0.7287	0.3191	0.8786
6	0.4488	0.8105	0.6237	0.3178	0.8191
7	0.1430	0.4130	0.2510	0.1518	0.4124

Table 3.3: Values of ϕ_{MIN} for different cases of shape mismatch.

Examination of Table 3.3 reveals that, unless the shape mismatch is fairly severe (i.e., the assumed λ is off by a factor 4 from the actual λ), the oversampled LSM-Gabor detector is most robust. Notice that for all cases (1, 3, 4, and 5) where the actual value of λ is 1, the values of ϕ_{MIN} for the SSTFT detector with rectangular window remain the same. This occurs because the rectangular window is independent of the shape parameter λ . Thus, the values of ϕ_{MIN} for the \mathbf{S}_r column reflect how well the 1-second rectangular window matches the one-sided exponential window with the given value of λ . It should also be noted that oversampling the LSM-Gabor detector does not provide a substantial improvement in performance over that of the critically sampled LSM-Gabor detector when there is shape mismatch. This is to be expected, as oversampling does not affect the shape of the window used in the Gabor transform in any way.

3.6.4 Effects of varying the number K of signal components

One other question to examine is how the performance of the five GLRT detectors is affected when the number of known signal components is reduced from $K = 6$ to $K = 2$. Table 3.4 below shows the values of ϕ_{MIN} for the five detectors for the same cases examined in Tables 3.2 and 3.3. The eleven cases listed in Table 3.4 are identified as follows:

1. arrival-time mismatch of 0
2. arrival-time mismatch of 0.25
3. arrival-time mismatch of 0.5
4. frequency mismatch of 0
5. frequency mismatch of 0.25
6. frequency mismatch of 0.5
7. shape mismatch with $\lambda_{true} = 1$ and $\lambda_{assumed} = 1$
8. shape mismatch with $\lambda_{true} = 0.5$ and $\lambda_{assumed} = 1$
9. shape mismatch with $\lambda_{true} = 1$ and $\lambda_{assumed} = 2$
10. shape mismatch with $\lambda_{true} = 1$ and $\lambda_{assumed} = 4$
11. shape mismatch with $\lambda_{true} = 1$ and $\lambda_{assumed} = 0.5$
12. shape mismatch with $\lambda_{true} = 2$ and $\lambda_{assumed} = 1$
13. shape mismatch with $\lambda_{true} = 4$ and $\lambda_{assumed} = 1$.

An added advantage to setting $K = 2$ is that the location minimization to determine location-independent values of ϕ_{MIN} can be done through exhaustive search, as there are only $\binom{16}{2} = 120$ location combinations when $K = 2$.

Comparison of Table 3.4 with Tables 3.2 and 3.3 reveals that all of the GLRT detectors perform better (i.e., have higher values of ϕ_{MIN}) when there are 2 signal components instead of 6. The reason for this is fairly intuitive: the signal energy $E = \vec{x}^H \vec{x}$ of the signal \vec{x} remains the same when $K = 2$, but the energy is distributed over a fewer number of signal components. Thus, the signal “stands out” more from the Gaussian noise, even when there is information mismatch. It is also interesting to

Case	Detector				
	Γ	B_c	S_r	S_c	B_o
1	0.7826	1	0.8851	0.5991	1
2	0.4692	0.4407	0.5791	0.4887	0.6125
3	0.2676	0.2689	0.3369	0.2894	1
4	0.7826	1	0.7683	0.5619	1
5	0.5429	0.6186	0.6471	0.5314	0.8304
6	0.2532	0.2886	0.3877	0.3253	1
7	0.7826	1	0.7607	0.4784	1
8	0.7815	0.8705	0.5934	0.3901	0.8705
9	0.5424	0.8580	0.7607	0.3914	0.8580
10	0.2325	0.5252	0.7607	0.1873	0.5252
11	0.7815	0.8705	0.7607	0.4357	0.8705
12	0.5424	0.8580	0.6795	0.4046	0.8580
13	0.2325	0.5252	0.3442	0.2047	0.5252

Table 3.4: Values of ϕ_{MIN} for different cases of location mismatch and shape mismatch, assuming $K = 2$ signal components.

note that the relative performances of the five detectors in Table 3.4 are exactly the same as the relative performances of the five detectors in Tables 3.2 and 3.3. Most importantly, this means that the oversampled LSM-Gabor detector is again the most robust among all the detectors in almost every case of information mismatch when $K = 2$.

3.6.5 Oversampling the other linear transforms

Given the dramatic improvement in GLRT detector robustness resulting from oversampling the LSM-Gabor transform, it natural to wonder whether oversampling would improve GLRT robustness for any of the other three critically sampled detectors examined in this section.

Oversampling the DBFM-Gabor transform would almost certainly improve GLRT robustness in the same way that oversampling the LSM-Gabor transform does. However, calculating the oversampled DBFM-Gabor transform coefficients requires the calculation of the continuous time biorthogonal function for the one-sided exponential window using an oversampled version of the Zak transform. As yet, no work

in this area has been published, although manuscripts are in preparation [FZ]. It should be noted that an extensive amount of research has recently been published regarding the oversampled Gabor transform calculated via the PBFM (periodized biorthogonal function method), and this work will be discussed in Chapter 4.

Oversampling the two SSTFT transforms turns out not to be advantageous. The reason for this is that the SSTFT is calculated using inner products. Because of this, the values of the SSTFT transform coefficients at the critically sampled lattice points remain unchanged at the corresponding oversampled lattice points. For every critically sampled lattice point, however, the oversampled SSTFT is evaluated at four neighboring lattice points on the oversampled lattice, and these neighboring transform coefficients are likely to have values close to the value of the transform coefficient at the critically sampled lattice point. The effect of this is that component localization is actually decreased when the SSTFT is oversampled.

The above property of the oversampled SSTFT is best illustrated by an example. Consider a signal with one component at time-frequency coordinate $(t, f) = (1, 1)$, with one-sided exponential window having parameter $\lambda = 1$. The magnitudes of the (1-second) rectangular-window SSTFT coefficients for this signal in the critically sampled and oversampled cases (along with the respective contour plots) are shown in Figure 3.9.

As noted in Chapter 2, the SSTFT is a blurred version of the Gabor transform, and this is already noticeable in the critically sampled case (left-hand side of Figure 3.9). However, the oversampled SSTFT is even more “blurred” than the critically sampled SSTFT. Figure 3.10 shows the corresponding plots for the SSTFT calculated with the one-sided exponential window (and no information mismatch). Again, the oversampled SSTFT is more blurred than the critically sampled SSTFT.

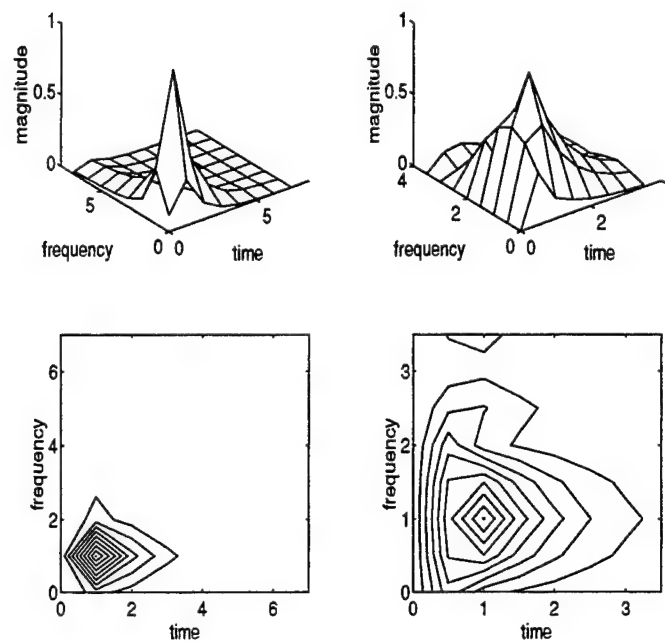


Figure 3.9: Left-hand side: magnitudes of the critically sampled SSTFT coefficients (rectangular window) and corresponding contour plot for one-component signal located at (1,1). Right-hand side: magnitudes of the oversampled, SSTFT coefficients (rectangular window) and corresponding contour plot for the same signal.

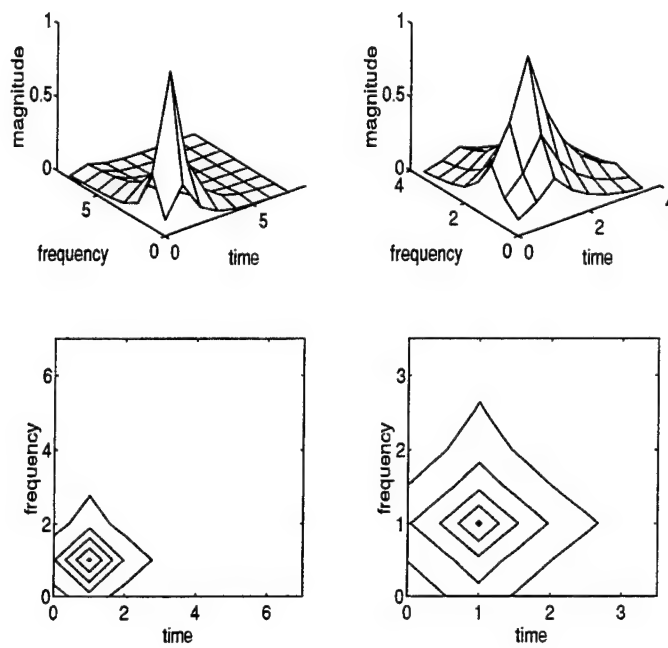


Figure 3.10: Left-hand side: magnitudes of the critically sampled SSTFT coefficients (one-sided exponential window) and corresponding contour plot for one-component signal located at (1,1). Right-hand side: magnitudes of the oversampled SSTFT coefficients (one-sided exponential window) and corresponding contour plot for the same signal.

3.7 Detection of Unknown-Location Signal Components in the PF Framework

To this point, the analysis of this chapter has focused on transient signal detection within the PF framework of signals whose components have *known locations* in the time-frequency plane. In the “known-location” case, knowledge of signal component locations is used to rearrange and partition the basic model (3.19) so that the GLRT detector can be designed (as in (3.24)) to test whether a portion of the transform coefficients (represented by \vec{c}_1) are zero or not.

This section examines transient signal detection within the PF framework of signals whose components have *unknown* locations. In fact, the *number* of signal components is also assumed to be unknown. The GLRT detector derived in this section is based on the oversampled LSM-Gabor transform, since the latter provided the most robust GLRT detector in the previous section in almost every case of information mismatch. Recall that for the oversampled Gabor transform, the transform matrix is

$$\mathbf{R} = \mathbf{B}_o = (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H, \quad (3.40)$$

where \mathbf{G}_o contains oversampled Gabor logons of the form

$$g_{mn}[k] = g_{mn}(k\Delta) = g(k\Delta - 0.5n)e^{j2\pi(0.5m)(k\Delta - 0.5n)}. \quad (3.41)$$

Rewriting the transform step (3.19) of the PF framework using $\mathbf{R} = \mathbf{B}_o$ gives

$$\vec{z} = \mathbf{B}_o \mathbf{G}_o \vec{c} + \mathbf{B}_o \vec{e} + \vec{w}, \quad (3.42)$$

where $\vec{w} \sim \mathcal{N}(\vec{0}, 2\mathbf{B}_o \mathbf{B}_o^H)$ and where \vec{e} is such that $\vec{e}^H \mathbf{G}_o = \vec{0}$. Note, however, that $\vec{e}^H \mathbf{G}_o = \vec{0}$ implies that $\mathbf{G}_o^H \vec{e} = \vec{0}$, so

$$\begin{aligned} \mathbf{B}_o \vec{e} &= (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H \vec{e} \\ &= (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \vec{0} \\ &= \vec{0}. \end{aligned} \quad (3.43)$$

Thus, the model mismatch vector \vec{e} can be disregarded in (3.42) above, giving

$$\vec{z} = \mathbf{B}_o \mathbf{G}_o \vec{c} + \vec{w}. \quad (3.44)$$

Since $\vec{w} \sim \mathcal{N}(\vec{0}, 2\mathbf{B}_o \mathbf{B}_o^H)$, the transformed observation $\vec{z} \sim \mathcal{N}(\mathbf{B}_o \mathbf{G}_o \vec{c}, 2\mathbf{B}_o \mathbf{B}_o^H)$.

The reason that \vec{e} could not be similarly disregarded for calculations involving the LSM-Gabor transform in the previous analysis involving known signal component locations is that the orthogonality constraint given in (3.31) only assumed

$$\vec{e}^H \mathbf{G}_{o_i} = \vec{0},$$

where the logons in \mathbf{G}_{o_1} corresponded to the known signal component locations of \vec{c}_1 . With unknown signal component locations, no such partitioning of \mathbf{G}_o into \mathbf{G}_{o_1} and \mathbf{G}_{o_2} is possible, so (3.43) applies and (3.44) is the model of interest. It should be noted that while model mismatch is disregarded in this section, information mismatch will play an important role in the analysis of the GLRT detector for unknown signal component locations (the “unknown location” or “UL” detector), as compared to the GLRT detector for known signal component locations (the “known location” or “KL” detector).

When signal component locations are unknown, the hypothesis test for (3.44) is simply

$$\begin{aligned} H_0 : \vec{c} &= \vec{0} \\ &vs. \\ H_1 : \vec{c} &\neq \vec{0}. \end{aligned} \quad (3.45)$$

For this hypothesis test, the logarithmic form of the GLRT is given by

$$t_{UL} = \max_{\vec{c}} \{4 \log p_1(\vec{z} | \vec{c})\} - 4 \log p_0(\vec{z}). \quad (3.46)$$

Now, assuming that \mathbf{B}_o is a left-inverse of \mathbf{G}_o (which is true if there is no information mismatch),

$$2 \log p_1(\vec{z}) = -MN \log 2\pi - \log |\Lambda| - (\vec{z} - \vec{c})^H \Lambda^{-1} (\vec{z} - \vec{c}), \quad (3.47)$$

where

$$\Lambda = 2\mathbf{B}_o \mathbf{B}_o^H = 2(\mathbf{G}_o^H \mathbf{G}_o)^{-1}. \quad (3.48)$$

The corresponding expression for $2 \log p_0(\vec{z})$ is obtained by substituting $\vec{c} = \vec{0}$ into (3.47). It is easily seen that $2 \log p_1(\vec{z})$ is maximized over \vec{c} when

$$\vec{c} = \vec{z}.$$

This is to be expected, as \vec{z} is a maximum likelihood estimate of \vec{c} [BS90]. Thus, the GLRT statistic in (3.46) simplifies to

$$t_{UL} = 2\vec{z}^H \Lambda^{-1} \vec{z} \quad (3.49)$$

$$= \vec{z}^H (\mathbf{B}_o \mathbf{B}_o^H)^{-1} \vec{z} \quad (3.50)$$

$$= \vec{z}^H (\mathbf{G}_o^H \mathbf{G}_o) \vec{z}. \quad (3.51)$$

To determine the distribution of t , note that since \vec{z} is MN -variate complex-Gaussian with mean $\mathbf{B}_o \mathbf{G}_o \vec{c}$ and covariance matrix $2\mathbf{B}_o \mathbf{B}_o^H$, the vector

$$\vec{\zeta} = (\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}} \vec{z} \quad (3.52)$$

is also MN -variate complex-Gaussian, but with mean $(\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}} \mathbf{B}_o \mathbf{G}_o \vec{c}$ and twice the *identity* covariance matrix. In a sense, $\vec{\zeta}$ represents the “whitened” transform coefficients. Examination of the GLRT statistic (3.51) above reveals that t_{UL} can be rewritten as

$$t_{UL} = \vec{\zeta}^H \vec{\zeta}, \quad (3.53)$$

which means that t_{UL} is χ^2 -distributed with $2MN$ degrees of freedom and noncentrality parameter

$$\nu = (\mathbf{B}_o \mathbf{G}_o \vec{c})^H (\mathbf{B}_o \mathbf{B}_o^H)^{-1} (\mathbf{B}_o \mathbf{G}_o \vec{c}). \quad (3.54)$$

One interpretation of (3.53) is that t_{UL} is simply the energy of the “whitened” transform coefficients.

Substituting $\vec{x} = \mathbf{G}_o \vec{c}$ into (3.54) reveals the dependence of the noncentrality parameter ν on \vec{x} :

$$\nu = \vec{x}^H \mathbf{B}_o^H (\mathbf{B}_o \mathbf{B}_o^H)^{-1} \mathbf{B}_o \vec{x}. \quad (3.55)$$

The expression for ν in (3.55) can be simplified again using the definition of \mathbf{B}_o ($\mathbf{B}_o = (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H$):

$$\nu = \vec{x}^H \mathbf{G}_o \mathbf{B}_o \vec{x} \quad (3.56)$$

$$= \vec{x}^H \mathbf{G}_o (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H \vec{x} \quad (3.57)$$

$$= \vec{x}^H \mathbf{P}_{\mathbf{G}_o} \vec{x}, \quad (3.58)$$

where $\mathbf{P}_{\mathbf{G}_o}$ is the projection of \mathbf{G}_o .

When it is in fact the case that $\vec{x} = \mathbf{G}_o \vec{c}$ for some \vec{c} , the expression (3.57) can be simplified even further:

$$\nu = \vec{x}^H \mathbf{G}_o (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H \vec{x} \quad (3.59)$$

$$= \vec{c}^H \mathbf{G}_o^H \mathbf{G}_o (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H \mathbf{G}_o \vec{c} \quad (3.60)$$

$$= \vec{c}^H \mathbf{G}_o^H \mathbf{G}_o \vec{c} \quad (3.61)$$

Thus, when there is no information mismatch, the noncentrality parameter of t_{UL} is simply the energy E of the signal \vec{x} .

3.8 Analysis of Unknown-Location (UL) GLRT Detector: Comparison With Known-Location (KL) Detector

This section analyzes the performance of the UL-detector t_{UL} given by (3.51) for the hypothesis test (3.45) and compares the performance of the UL-detector with that of the KL-detector t (which will be denoted t_{KL} in this section) given by (3.25) for some sample signals. Both the UL-detector and the KL-detector are GLRT statistics based on the oversampled LSM-Gabor transform ($\alpha = \beta = 0.5$). For convenience, the subscript ‘o’ (for “oversampled”) will be dropped from \mathbf{G}_o in this section, but $\mathbf{G} = \mathbf{G}_o$ will be implied throughout. Similarly, the oversampled LSM-Gabor transform matrix \mathbf{B}_o will be denoted by \mathbf{B} .

Recall that the KL-detector assumes the signal \vec{x} is of the form

$$\vec{x} = \mathbf{G}_1 \vec{c}_1, \quad (3.62)$$

with

$$\vec{c} = \begin{bmatrix} \vec{c}_1 \\ \vec{0} \end{bmatrix}. \quad (3.63)$$

In [PF92], Porat and Friedlander showed that when the signal \vec{x} is of the form $\vec{x} = \mathbf{G}\vec{c}$ (but not necessarily of the form shown in (3.62)), the noncentrality parameter ν_{KL} of the KL-detector t_{KL} is given by

$$\nu_{KL} = \vec{c}_1^H \mathbf{G}_1^H \mathbf{G}_1 \vec{c}_1, \quad (3.64)$$

which equals $\vec{x}^H \vec{x}$ if \vec{x} is of the form assumed by (3.62). Thus, when the information about the locations of the signal components is correct, the noncentrality parameters of the KL-detector and the UL-detector are exactly the same, both equalling the energy E of the signal \vec{x} . However, recall that the χ^2 -distribution of the KL-detector has only $2K$ degrees of freedom (where K is the number of signal components), whereas the χ^2 -distribution of the UL-detector has $2MN$ degrees of freedom. Since $K \ll MN$ by the sparseness assumption A2, this means that when there is no information mismatch, the KL-detector will perform considerably better than the UL-detector⁷.

However, note that when it is *not* the case that \vec{x} is of the form $\vec{x} = \mathbf{G}_1 \vec{c}_1$, the noncentrality parameter ν_{KL} will *not* equal the energy $E = \vec{x}^H \vec{x}$ of \vec{x} , even if \vec{x} is of

⁷For a given probability of false alarm and a given noncentrality parameter under H_1 , the detection probability using either of the χ^2 statistics is a monotonically *decreasing* function of the number of degrees of freedom.

the form $\vec{x} = G\vec{c}$. In other words, if \vec{x} can be represented as $\vec{x} = G\vec{c}$, with

$$\vec{c} = \begin{bmatrix} \vec{c}_1 \\ \vec{c}_2 \end{bmatrix} \quad (3.65)$$

and $\vec{c}_2 \neq \vec{0}$, the value of ν_{KL} will only be a fraction of the energy E of \vec{x} . On the other hand ν_{UL} , as shown in (3.61), remains equal to E regardless of the locations of the signal components (so long as \vec{x} can be expressed as $\vec{x} = G\vec{c}$). In this case, there is a tradeoff for the two detectors between the fewer degrees of freedom in the KL-detector and the larger noncentrality parameter of the UL-detector. It is to be expected, then, that the relative performances of the two detectors depends on how accurate the assumptions are concerning the signal component locations.

3.8.1 Simulation examples



It would be beneficial at this point to consider some specific examples. As before, the underlying parameters of the model (see (3.38)) are assumed to be $M = N = 8$, $L = 256$, and $\Delta = \frac{1}{32}$. First, consider the signal \vec{x} whose samples are given by


$$x(k\Delta) = \sqrt{2}e^{-(k\Delta-1.5)}e^{j3\pi(k\Delta-1.5)}u(k\Delta-1.5), \quad k = 0, 1, \dots, L-1, \quad (3.66)$$

a one-component signal with one sided-exponential window and arrival time 1.5, modulation frequency 1.5, and decay parameter $\lambda = 1$. The KL detector assumes that the arrival time and modulation frequency of the one-component signal are indeed 1.5. If this is true, then the KL detection statistic is noncentral χ^2 -distributed under hypothesis H_1 with 2 degrees of freedom and noncentrality parameter $\nu_{KL} = 33.0103$. Under hypothesis H_0 , the KL detection statistic is central χ^2 -distributed with 2 degrees of freedom. The respective probability distributions are shown in Figure 3.11.

The UL detector makes no assumptions on the arrival time and modulation frequency of the one-component signal. Thus, the UL detection statistic under hypothesis H_1 is noncentral χ^2 -distributed with 128 degrees of freedom and noncentrality parameter $\nu_{UL} = 33.0103$, while under hypothesis H_0 it is central χ^2 -distributed with 128 degrees of freedom. The respective probability distributions are shown in Figure 3.12.

Figures 3.11 and 3.12 indicate that when there is no mismatch, the KL detector should perform better. This is verified in Figure 3.13, where the probability of detection P_d for the KL and UL detectors are plotted as a function of the signal-to-noise ratio for a false alarm probability $P_{fa} = 0.01$. The SNR was varied by changing the magnitude of the signal \vec{x} .

Figure 3.11: χ^2 distributions, 2 degrees of freedomFigure 3.12: χ^2 distributions, 128 degrees of freedom

Figure 3.13: P_d vs. SNR for $P_{fa} = 0.01$

Suppose, however, that the assumed time-frequency location of the one-component signal is incorrect. This would mean that the arrival time αn of the component is not 1.5 or that the modulation frequency βm of the component is not 1.5 or that neither quantity is 1.5. Since the UL detector makes no assumptions on the location of the component, its performance will remain the same as long as the actual location of the component remains within the finite expected ranges of arrival time and modulation frequency (both 0 to 3.5 in this simulation example). However, the performance of the KL detector deteriorates the further the actual time-frequency coordinates of the signal are from the expected coordinates, since the KL detector makes use of the location information. Moreover, this deterioration is quite rapid. When the actual (t, f) coordinates of the component are simply $(2, 2)$ instead of $(1.5, 1.5)$, the noncentrality parameter of the KL statistic under hypothesis H_1 drops from 33.0103 all the way to 3.5051. As shown in Figure 3.14, there is much less separation now between the distributions of the KL detector under H_0 and H_1 , respectively.

Indeed, Figure 3.15 shows that the performance of the KL detector drops below that of the UL detector with this minimal amount of location mismatch.



Figure 3.14: χ^2 distributions, 2 degrees of freedom

Figure 3.15: P_d vs. SNR for $P_{fa} = 0.01$. In the mismatch case, the actual time-frequency coordinates of the signal component are (2, 2) instead of (1.5, 1.5). The UL curve is the same in either case.




Figure 3.16: P_d vs. SNR for $P_{fa} = 0.001$. In the mismatch case, the actual time-frequency coordinates of the signal component are $(2, 2)$ instead of $(1.5, 1.5)$. The UL curve is the same in either case.

It should be clarified at this point that location mismatch for the KL detector means that the actual t - f coordinates of the signal component differ from the expected coordinates $(1.5, 1.5)$. On the other hand, location mismatch for the UL detector signifies that the actual t - f coordinates of the signal component do not lie on the oversampled Gabor lattice.

Figure 3.16 shows the corresponding performance curves when P_{fa} is set to 0.001 instead of 0.01 as in Figure 3.15. Again, the performance of the KL detector drops below that of the UL detector when the actual time-frequency coordinates of the signal component are simply $(2, 2)$ instead of $(1.5, 1.5)$.

Figure 3.17 displays the ROCs, P_d versus P_{fa} , for the two detectors when the SNR is set to approximately -9.0 dB. The detection probability of the KL detector when there is no mismatch is essentially 1 except at the very smallest false alarm probabilities. However, the ROCs curve for the KL detector drops dramatically with the same location mismatch as above.




Figure 3.17: ROCs: P_d vs. P_{fa} for $\text{SNR} \approx -9.0$ dB. In the mismatch case, the actual time-frequency coordinates of the signal component are $(2, 2)$ instead of $(1.5, 1.5)$. The UL curve is the same in either case.

A broader indication of the sensitivity of the KL detector to location mismatch is given by Figure 3.18, which shows the dropoff in the detection probability P_d for the KL detector with respect to P_d for the UL detector as the actual time-frequency location of the signal component varies from the expected $(1.5, 1.5)$. It can be seen from Figure 3.18 that the performance of the KL detector drops below that of the UL detector unless the actual time-frequency coordinates of the signal component are *exactly* as expected or unless one of the coordinates is exact and the other is off by 0.5.

In Figure 3.18, the values of P_d for the UL detector remain constant as long as the actual time-frequency coordinates of the signal component lie somewhere on the oversampled Gabor lattice (i.e., (t, f) is such that $t = 0.5n$ and $f = 0.5m$, with $m, n \in \{0, 1, \dots, 7\}$). The performance of the UL detector does suffer a dropoff when the actual time-frequency coordinates of the signal component do *not* lie exactly on the oversampled Gabor lattice. However, this dropoff is not nearly as severe as that suffered by the KL detector. Figure 3.19 shows the variation in P_d for both the KL and UL detectors over a smaller portion of the time-frequency plane, when the expected time-frequency location of the signal component is $(1, 1)$. Though P_d for the UL detector is no longer constant in Figure 3.19, the variation is minimal when compared to P_d for the KL detector.




Figure 3.18: Variation of P_d for the KL detector as the actual time-frequency coordinates of the signal component vary from the expected (1.5, 1.5). The corresponding P_d for the UL detector is shaded. $\text{SNR} \approx -9.0$ dB, $P_{fa} = 0.01$.




Figure 3.19: Variation of P_d for the KL detector as the actual time-frequency coordinates of the signal component vary from the expected $(1, 1)$. The corresponding P_d for the UL detector is shaded. $\text{SNR} \approx -9.0$ dB, $P_{fa} = 0.01$.




Figure 3.20: P_d vs. P_{fa} for $\text{SNR} \approx -9.0$ dB. In the mismatch case, the actual time-frequency coordinates of the signal component are (1.7, 1.7) instead of (1.5, 1.5).

Figure 3.20 shows the dropoff in the ROCs of the UL detector as the signal component's t - f coordinates vary from a no mismatch case (coordinates somewhere on the oversampled Gabor lattice) to a mismatch case (coordinates of (1.7, 1.7)). It is clear that the dropoff is not significant.

Careful examination of the data used to generate Figure 3.19 reveals that, for a given component arrival time, P_d for the UL detector does not vary at all with component frequency, even if there is location mismatch. With this in mind, the variation of P_d for the UL detector as the signal component's actual *arrival time* varies (while the *frequency* is fixed) is shown in Figure 3.21, which is a cross-section of Figure 3.19 at $f = 1$.

The overall conclusion from all the simulation examples is that the UL detector, as one would expect, is more robust than the KL detector when there is location mismatch. One should only use the KL detector when one is very confident about the time-frequency locations of the signal components. From Figures 3.18 and 3.19, "very confident" can be taken to mean that the signal component locations are either known exactly or known to within adjacent⁸ coordinates on the oversampled Gabor lattice. Although the simulation examples above used one-component signals, the above conclusion is obviously applicable to multiple-component signals.

⁸"Adjacent" here means horizontally or vertically adjacent, not diagonally adjacent.




Figure 3.21: P_d vs. time for $\text{SNR} \approx -9.0$ dB, $f = 1$.

3.8.2 Comparison: UL detector and energy detector

One final question to be addressed in this section is why the UL detector provides any advantage over a simple energy detector when the component locations are unknown. This is an especially relevant question considering that the UL detection statistic (3.51) can be interpreted as the energy of the “whitened” transform coefficients.

To answer this question, consider again the PF framework representation of the observation vector \vec{y} :

$$\vec{y} = \vec{x} + \vec{v} = \mathbf{G}\vec{c} + \vec{v}. \quad (3.67)$$

The energy detector for (3.67) is simply $t_E = \vec{y}^H \vec{y}$. Since $\vec{y} \sim \mathcal{N}(\vec{x}, 2\mathbf{I})$, it is clear that, for the hypothesis test (3.45) above, t_E is central χ^2 -distributed with $2L$ degrees of freedom under H_0 and noncentral χ^2 -distributed with $2L$ degrees of freedom and noncentrality parameter $\vec{x}^H \vec{x} = E$ under H_1 .

Thus, the distributions t_E for the hypothesis test (3.45) are essentially the same as the distributions of t_{UL} given above, *except* that the χ^2 -distributions for t_E have $2L$ degrees of freedom, while the χ^2 -distributions for t_{UL} have $2MN$ degrees of freedom. Now the PF framework assumes (see assumption A1 in section 3.2) that $MN < L$. Indeed, in the simulation examples of this section, $MN = 64 \ll L = 256$. Thus, because of the nature of the χ^2 distribution, the statistic t_{UL} will always perform better than the statistic t_E , unless the PF model (3.2) for the signal \vec{x} is grossly inaccurate.

Essentially, one can conclude from the above that if an L -dimensional signal is known to belong to an MN -dimensional linear subspace, it is beneficial to perform a linear transform on the observation to reduce the detection problem from an L -dimensional one to an MN -dimensional one, *provided* one has a reasonable knowledge of the linear subspace.

3.9 Scharf-Friedlander Framework: Matched Subspace Detectors

Recently, Scharf and Friedlander [SF94] developed a generalization of the PF framework. The Scharf-Friedlander framework – which will be referred to here, naturally, as the SF framework – includes *subspace interferences* as part of its model. In other words, the observation \vec{y} is again expressed as

$$\vec{y} = \vec{x} + \vec{n} = \mathbf{G}\vec{c} + \vec{n}, \quad (3.68)$$

except that the noise \vec{n} has distribution $\vec{n} \sim \mathcal{N}(\mathbf{Q}\vec{\sigma}, \mathbf{I})$, where $\mathbf{Q}\vec{\sigma}$ represents the subspace interference⁹. $\vec{\sigma}$ does not necessarily have to have the same dimensions as \vec{c} ; thus, the interference matrix \mathbf{Q} is assumed to be $(L \times s)$, while $\vec{\sigma}$ is assumed to be $(s \times 1)$.

Equivalently, the observation \vec{y} can be expressed as

$$\vec{y} = \mathbf{G}\vec{c} + \mathbf{Q}\vec{\sigma} + \vec{v}, \quad (3.69)$$

where $\vec{v} \sim \mathcal{N}(\vec{0}, \mathbf{I})$ ¹⁰.

Using the observation model given by (3.69), Scharf and Friedlander derived the GLRT statistic for the hypothesis test that was given in (3.45) above:

$$\begin{aligned} \mathbf{H}_0 : \vec{c} &= \vec{0} \\ &vs. \\ \mathbf{H}_1 : \vec{c} &\neq \vec{0}. \end{aligned} \quad (3.70)$$

⁹Scharf and Friedlander assumed real-valued signals and real-valued noise. The extension of their analysis to complex-valued signals would involve similar steps to those taken in previous sections of this chapter.

¹⁰Scharf and Friedlander also generalized their framework to allow the covariance matrix of \vec{v} to be an arbitrary multiple of the identity matrix, $\sigma^2 \mathbf{I}$. This generalization is a simple one, and it is assumed here that $\sigma^2 = 1$.

The GLRT statistic derived by Scharf and Friedlander is given by

$$t_{SF} = \vec{y}^H \mathbf{P}_Q^\perp \mathbf{E}_{GQ} \mathbf{P}_Q^\perp \vec{y}, \quad (3.71)$$

where

$$\mathbf{P}_Q^\perp = \mathbf{Q}^\perp ((\mathbf{Q}^\perp)^H \mathbf{Q}^\perp)^{-1} (\mathbf{Q}^\perp)^H$$

and

$$\mathbf{E}_{GQ} = \mathbf{G} (\mathbf{G}^H \mathbf{P}_Q^\perp \mathbf{G})^{-1} \mathbf{G}^H \mathbf{P}_Q^\perp.$$

Without exploring the intricacies of the detection statistic t_{SF} , it is interesting to note here that when $\mathbf{Q} = \mathbf{0}$, t_{SF} reduces to t_{UL} if the transform matrix \mathbf{R} in (3.19) is taken to be the pseudoinverse, $\mathbf{R} = \mathbf{B} = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H$. To see this, assume $\mathbf{Q} = \mathbf{0}$, so that $\mathbf{Q}^\perp = \mathbf{I}$. This makes

$$\mathbf{P}_Q^\perp = \mathbf{I}$$

and

$$\mathbf{E}_{GQ} = \mathbf{G} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H,$$

and the statistic t_{SF} simplifies as follows:

$$t_{SF} = \vec{y}^H \mathbf{P}_Q^\perp \mathbf{E}_{GQ} \mathbf{P}_Q^\perp \vec{y} \quad (3.72)$$

$$= \vec{y}^H \mathbf{G} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \vec{y} \quad (3.73)$$

$$= \vec{y}^H \mathbf{G} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \mathbf{G} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \vec{y} \quad (3.74)$$

$$= \vec{z}^H \mathbf{G}^H \mathbf{G} \vec{z}, \quad (3.75)$$

provided $\vec{z} = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \vec{y}$. Note that the form of t_{SF} in (3.75) is identical to the form of t_{UL} in (3.51).

Thus, the analysis of the UL detector in sections 3.7–3.8 is a special case of the SF framework. It should be noted that Scharf and Friedlander in [SF94] did not attempt to apply linear TFRs to signal detection problems within their framework; in fact, they based their GLRT statistic on the observation \vec{y} , not on any transformed observation $\vec{z} = \mathbf{R}\vec{y}$. Thus, the analysis of sections 3.7–3.8, while a special case of the SF framework, remains distinct from the work of Scharf and Friedlander in that the latter does not approach the signal detection problem from the point of view of linear TFRs.

It should be noted further that Friedlander and Porat in [FP92] derived t_{UL} as the expression in (3.49), but they focused immediately thereafter on the KL case, deriving t_{KL} without analyzing t_{UL} . Again, the form of t_{UL} was derived in this paper, but the analysis of sections 3.7–3.8 remains distinct from the work of Friedlander and Porat.

3.10 Conclusions

This chapter applied the discrete Gabor transform to the problem of detecting finite-duration, discrete-time signals. The first part of the chapter focused on the case where the time-frequency locations of the signal components were assumed *known*. In a comparison of GLRT detectors based on five linear TFRs for this case, it was found that the detector based on the oversampled LSM-Gabor transform was generally most robust in the presence of signal mismatch, both information mismatch and model mismatch.

The second part of the chapter focused on the case where the time-frequency locations of the signal components were assumed *unknown*. The GLRT detector based on the oversampled LSM-Gabor transform was derived for this case, and the performance of this “unknown location” (UL) detector was compared to the performance of the corresponding “known location” (KL) detector. It was found, as expected, that the UL detector was more robust in the presence of location mismatch.

Two possible topics for further investigation are the application of oversampling to the DBFM-Gabor transform and the analysis of transform-based detection within the Scharf-Friedlander (SF) framework of matched subspace detection. Both of these issues are being researched currently [FZ, SF94].

3.A Appendix: Restrictions on ν_{min} in Section 3.4

This appendix confirms the necessity of the condition $\rho \leq \phi$ for the minimum solution $\nu_{min} = E \left(\sqrt{(1-\rho)\phi} - \sqrt{(1-\phi)\rho} \right)^2$ in section 3.4.

First, recall that the original minimization problem is to find

$$\min_{\vec{e}} \nu = \min_{\vec{e}} (\mathbf{G}_1 \vec{c}_1 + \vec{e})^H \mathbf{P} (\mathbf{G}_1 \vec{c}_1 + \vec{e}), \quad (3A.1)$$

subject to the conditions

$$\begin{aligned} \vec{e}^H \vec{x}_0 &= \vec{0} \\ \vec{e}^H \vec{e} &= E - \vec{x}_0^H \vec{x}_0, \end{aligned} \quad (3A.2)$$

where $\vec{x}_0 = \mathbf{G}_1 \vec{c}_1$.

This is a *constrained minimization* problem. Letting $f(\vec{e})$ be the expression for ν_{min} on the right-hand side of (3A.1), and letting $h(\vec{e})$ be the vector-valued function

given by

$$h_1(\vec{e}) = \vec{e}^H \vec{e} - E + \vec{x}_0^H \vec{x}_0 \quad (3A.3)$$

$$h_2(\vec{e}) = 2\vec{e}^H \vec{x}_0, \quad (3A.4)$$

the first-order necessary condition to find a vector \vec{e}^* that minimizes $f(\vec{e})$ is that $\exists \vec{\theta}$ such that

$$\nabla_{\vec{e}} f(\vec{e}^*) + \vec{\theta}^H \nabla_{\vec{e}} \vec{h}(\vec{e}^*) = \vec{0}, \quad (3A.5)$$

where $\vec{\theta} = \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}$. The first-order condition (3A.5) can then be combined with the two constraints in (3A.2) to solve for \vec{e}^* and $\vec{\theta}$.

In [PF92], Porat and Friedlander performed the above calculations by expressing \vec{e}^* in terms of θ_1 and θ_2 , then expressing θ_2 in terms of θ_1 , and, finally, solving for θ_1 . They found two solutions for θ_1 , given by

$$\theta_1^+ = -1 + \phi + \sqrt{\frac{\phi}{\rho}(1-\phi)(1-\rho)} \quad (3A.6)$$

$$\theta_1^- = -1 + \phi - \sqrt{\frac{\phi}{\rho}(1-\phi)(1-\rho)}. \quad (3A.7)$$

Porat and Friedlander assumed that θ_1^+ is the solution that generates ν_{min} given in (3.33); however, they did not check the second order conditions that determine *when* θ_1^+ corresponds to a minimum. This is done below.

The second-order necessary condition for a vector \vec{e}^* to minimize $f(\vec{e})$ is that $\exists \vec{\theta}$ such that

$$\nabla_{\vec{e}}^2 f(\vec{e}^*) + \vec{\theta}^H \nabla_{\vec{e}}^2 \vec{h}(\vec{e}^*) \quad (3A.8)$$

is positive semidefinite [Lue84]. In (3A.8), $\vec{\theta}^H \nabla_{\vec{e}}^2 \vec{h}(\vec{e}^*)$ is defined as

$$\vec{\theta}^H \nabla_{\vec{e}}^2 \vec{h}(\vec{e}^*) = \theta_1 \nabla_{\vec{e}}^2 h_1(\vec{e}^*) + \theta_2 \nabla_{\vec{e}}^2 h_2(\vec{e}^*).$$

After some calculations, the matrix in (3A.8) for the above constrained minimization problem is reduced to

$$2(\mathbf{P} + \theta_1 \mathbf{I}), \quad (3A.9)$$

Assuming that \mathbf{P} is full rank and has eigenvalue decomposition $\mathbf{P} = \mathbf{Q}^H \mathbf{D} \mathbf{Q}$, the matrix in (3A.9) can be rewritten

$$2\mathbf{Q}^H (\mathbf{D} + \theta_1 \mathbf{I}) \mathbf{Q}. \quad (3A.10)$$

Since \mathbf{P} is a projection matrix, its eigenvalues are real and positive (or 0). Assuming that the eigenvalues of \mathbf{P} can occur arbitrarily close to 0, it is clear that the matrix (3A.10) is positive semidefinite only if $\theta_1 \geq 0$.

Since $\phi \leq 1$, it is clear that $\theta_1^- \geq 0$ only in the degenerate case¹¹ where $\phi = 1$. The condition for $\theta_1^+ \geq 0$ is derived as follows (ignoring the degenerate case $\phi = 1$):

$$\begin{aligned}
 & \Rightarrow -1 + \phi + \frac{\theta_1^+}{\sqrt{\frac{\phi}{\rho}(1-\phi)(1-\rho)}} \geq 0 \\
 & \Rightarrow \sqrt{\frac{\phi}{\rho}(1-\phi)(1-\rho)} \geq 1 - \phi \\
 & \Rightarrow \frac{\phi}{\rho}(1-\phi)(1-\rho) \geq (1-\phi)^2 \\
 & \Rightarrow \phi(1-\rho) \geq \rho(1-\phi) \\
 & \Rightarrow \frac{\phi(1-\rho)}{\rho} \leq \phi.
 \end{aligned} \tag{3A.11}$$

The inequality is preserved during the squaring operation of the fourth step because both sides of the inequality are positive.

Thus, the condition $\rho \leq \phi$ is necessary in order for the expression (3.33) for ν_{min} to be valid.

¹¹When $\phi = 1$, $\theta_1 = 0$ and $\nu_{min} = 0$. The expression for ν_{min} in (3.33) was derived assuming $\phi \neq 1$.

Improving the Oversampled Gabor Transform

4.1 Introduction

This chapter focuses on some of the factors that comprise an effective linear TFR: stability in the magnitudes of the transform coefficients (referred to here as *coefficient stability*), accuracy of the representation, and computational simplicity. Section 4.2 presents a comparison of coefficient stabilities for the five linear TFRs examined in the last chapter, and it is shown that the oversampled LSM-Gabor transform (abbreviated in this chapter as OS-LSM-GT) is extremely *unstable* in comparison with the other transforms. It is then explained how the GLRT detector based on the OS-LSM-GT was able to outperform all the others in the analysis of Chapter 3 despite this instability.

The remainder of the chapter is devoted to investigating methods of improving the coefficient stability of the OS-LSM-GT. Section 4.3 examines five methods designed to improve the stability of the OS-LSM-GT while remaining within the Porat-Friedlander (PF) framework:

- rank reduction (via singular value decomposition)
- regularization
- the principal components method [TKK82, TK82]
- solution with energy constraints
- solution with transform “whitening.”

Section 4.4 compares all five of the above methods with the original OS-LSM-GT by examining each method's transform coefficients for some simple signals. Three performance measures are used for the comparison, all of which measure in different ways the accuracies of the various methods.

Sections 4.5–4.7 investigate a different approach toward calculating the oversampled Gabor transform coefficients. Section 4.5 explains the framework of Wexler and Raz (to be called the WR framework), and it examines the assumptions of this framework in comparison to the assumptions of the PF framework. Section 4.6 presents several methods of calculating the oversampled Gabor transform coefficients within the WR framework. Section 4.7 compares the WR methods of calculating the oversampled Gabor transform to the methods of sections 4.5–4.7, using the same performance measures and signals as presented in section 4.4.

Finally, section 4.8 summarizes the conclusions of this chapter and presents some ideas for future work.

4.2 Stability Analysis of TFRs in the PF Framework

This section compares the stability of the five linear TFRs compared in section 3.6 of the last chapter. Within the PF framework, consider the representation of a transformed observation vector \vec{z} , derived from a signal \vec{x} in complex, white Gaussian noise \vec{v} , assuming a generic transform matrix \mathbf{R} :

$$\vec{z} = \mathbf{R}\vec{y} = \mathbf{R}\vec{x} + \mathbf{R}\vec{v}, \quad (4.1)$$

where $\vec{v} \sim \mathcal{N}(\vec{0}, 2\mathbf{I})$. It is easy to see that $\vec{z} \sim \mathcal{N}(\mathbf{R}\vec{x}, 2\mathbf{R}\mathbf{R}^H)$. Thus, the terms in the covariance matrix $2\mathbf{R}\mathbf{R}^H$ determine how much the magnitudes of the transform coefficients within the transform vector \vec{z} are likely to vary with different realizations of the noise. The latter characteristic will be referred to here as the *coefficient stability* of the transform. “Coefficient stability” will also be used to refer to the relative variation in the transform coefficient magnitudes with small changes in the signal.

The five linear TFRs examined in section 3.6 – the critically sampled DBFM-Gabor transform with one-sided exponential window (with transform matrix $\mathbf{\Gamma}$), the critically sampled LSM-Gabor transform with one-sided exponential window (\mathbf{B}_c), the critically sampled SSTFT with rectangular window (\mathbf{S}_r), the critically sampled SSTFT with one-sided exponential window (\mathbf{S}_e), and the oversampled LSM-Gabor

transform with one-sided exponential window (\mathbf{B}_e) – are examined for coefficient stability in Table 4.1 below, which presents eight measures of element size for the covariance matrices corresponding to the various transforms:

1. the condition number, the ratio of the largest singular value of the covariance matrix to the smallest
2. the 2-norm, the largest singular value of the covariance matrix
3. the 1-norm, the largest column sum of the covariance matrix
4. the ∞ -norm, the largest row sum of the covariance matrix
5. the Frobenius norm, the square root of the sum of the squares of all the elements in the covariance matrix
6. the maximum diagonal element of the covariance matrix, corresponding to the maximum variance of any one transform coefficient in \vec{z}
7. the average of the diagonal elements in the covariance matrix, corresponding to the average variance of the transform coefficients in \vec{z}
8. the minimum diagonal element of the covariance matrix, corresponding to the minimum variance of any one transform coefficient in \vec{z} .

With all of the above measures, smaller values indicate greater coefficient stability in the corresponding transform.

The transform matrices in Table 4.1 use the same parameters within the PF framework that were assumed in the last chapter: $M = N = 8$, $L = 256$, and $\Delta = \frac{1}{32}$ (see equation (3.2) to review the meanings of these parameters). All of the transform matrices were created to detect signal components with one-sided exponential windows having decay parameter $\lambda = 1$.

Examination of Table 4.1 reveals that the critically sampled transforms, comprising the first four columns of the table, are all fairly stable, with the transform based on \mathbf{S}_r (the SSTFT with rectangular window) an order of magnitude more stable than the others¹. However, the elements of the covariance matrix $2\mathbf{B}_o\mathbf{B}_o^H$ corresponding to the OS-LSM-GT have extremely large magnitudes, indicating that the oversampled LSM-Gabor transform is very *unstable*. Note that the critically sampled LSM-Gabor transform based on \mathbf{B}_e is relatively stable by comparison.

¹For the critically sampled RW-SSTFT, the covariance matrix $2\mathbf{S}_r\mathbf{S}_r^H$ is a multiple of the identity matrix.

Measure	Covariance Matrix				
	$2\Gamma\Gamma^H$	$2\mathbf{B}_c\mathbf{B}_c^H$	$2\mathbf{S}_r\mathbf{S}_r^H$	$2\mathbf{S}_e\mathbf{S}_e^H$	$2\mathbf{B}_o\mathbf{B}_o^H$
condition number	18.6279	18.6279	1	18.6279	2.209e10
2-norm	0.3292	0.3072	0.0625	0.2368	2.656e8
1-norm	0.4018	0.3798	0.0625	0.3034	4.704e8
∞ -norm	0.4018	0.3798	0.0625	0.3034	4.704e8
Frobenius norm	1.0456	0.9610	0.50	0.6306	5.820e8
max(diag(\cdot))	0.1098	0.1046	0.0625	0.0644	7.512e7
avg(diag(\cdot))	0.1098	0.0992	0.0625	0.0632	2.382e7
min(diag(\cdot))	0.0968	0.0814	0.0625	0.0558	2.926e4

Table 4.1: Stability of coefficient estimates, $M=N=8$, $L=256$, $\Delta = \frac{1}{32}$

The instability in the OS-LSM-GT occurs even if it is performed on a signal *without* noise, because the pseudoinverse

$$\mathbf{B}_o = (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H$$

contains the inverse $(\mathbf{G}_o^H \mathbf{G}_o)^{-1}$ by definition. Indeed, the covariance matrix $2\mathbf{B}_o\mathbf{B}_o^H$ simplifies to $2\mathbf{B}_o\mathbf{B}_o^H = 2(\mathbf{G}_o^H \mathbf{G}_o)^{-1}$, so the source of instability is the same whether there is Gaussian noise added or not. The following example illustrates the lack of coefficient stability in the OS-LSM-GT.

Consider a one-component signal with one-sided exponential window (decay parameter $\lambda = 1$) and (t, f) coordinates $(1, 1)$. The magnitudes of the OS-LSM-GT coefficients are shown in Figure 4.1. Notice that when there is no location mismatch, the coefficient at $(1, 1)$ is perfectly represented.




Figure 4.1: Magnitudes of oversampled LSM-Gabor coefficients of one-component signal, located at (1, 1).

Now consider the magnitudes of the OS-LSM-GT coefficients of the signal shown in Figure 4.2, where the time coordinate of the signal component has been changed very slightly from 1 to 1.05. Ideally, the transform coefficients should not change very much from those depicted in Figure 4.1. However, the vertical axis of Figure 4.2 indicates that the calculated transform coefficients are now in the thousands. This is a vivid illustration of the lack of coefficient stability in the OS-LSM-GT: a slight change in the parameters of the signal produces a drastic change in the transform coefficients². In fact, the OS-LSM-GT coefficients will only be on the same order of magnitude as the actual coefficients of the signal if there is no arrival-time mismatch whatsoever. It is very impractical to assume that the latter will be the case, so the need to modify the OS-LSM-GT to improve its coefficient stability is obvious.

It should be explained at this point why the GLRT detector based on the OS-LSM-GT proved to be more robust in the presence of mismatch than the other four detectors in the analysis of the Chapter 3, given the lack of coefficient stability of the OS-LSM-GT in the presence of mismatch. This is best explained using singular value decomposition.

²It should be noted that the instability of the OS-LSM-GT is not as severe when there is frequency mismatch. Thus, this and future examples in this dissertation focus on cases of arrival-time mismatch.




Figure 4.2: Magnitudes of oversampled LSM-Gabor coefficients of one-component signal, located at (1.05, 1).

The *singular value decomposition* (SVD) of an $n \times p$ matrix A is given by

$$A = U \Sigma V^H, \quad (4.2)$$

where U ($n \times p$) and V ($p \times p$) are unitary and where Σ is a diagonal matrix containing the *singular values* σ_i of A . When A has small singular values, the elements of A^{-1} will be large.

Suppose then that the SVD of G_o is given by

$$G_o = U_{G_o} \Sigma_{G_o} V_{G_o}^H. \quad (4.3)$$

Then B_o is given by

$$\begin{aligned} B_o &= (G_o^H G_o)^{-1} G_o^H \\ &= (V_{G_o} \Sigma_{G_o} U_{G_o}^H U_{G_o} \Sigma_{G_o} V_{G_o}^H)^{-1} V_{G_o} \Sigma_{G_o} U_{G_o}^H \\ &= (V_{G_o} \Sigma_{G_o}^2 V_{G_o}^H)^{-1} V_{G_o} \Sigma_{G_o} U_{G_o}^H \\ &= V_{G_o} \Sigma_{G_o}^{-2} V_{G_o}^H V_{G_o} \Sigma_{G_o} U_{G_o}^H \\ &= V_{G_o} \Sigma_{G_o}^{-1} U_{G_o}^H. \end{aligned} \quad (4.4)$$

Thus, the singular values of B_o are the inverses of the singular values of G_o . If G_o has any small singular values, then the corresponding singular values of B_o will be large. When G_o contains oversampled Gabor logons with one-sided exponential

windows (parameter $\lambda = 1$), the smallest singular value of \mathbf{G}_o is approximately 0.000087. This leads to large elements in \mathbf{B}_o , which produces the instability in the OS-LSM-GT noted above.

However, the GLRT statistic t in the PF framework is formed by taking the energy of the “whitened” transform coefficients $\vec{\zeta}$:

$$t = \vec{\zeta}^H \vec{\zeta},$$

where $\vec{\zeta} = (\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}} \mathbf{B}_o \vec{y}$. Consider, then, the SVD of the matrix $\mathbf{W} = (\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}} \mathbf{B}_o$ used to form $\vec{\zeta}$:

$$\begin{aligned} \mathbf{W} &= (\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}} \mathbf{B}_o \\ &= (\mathbf{G}_o^H \mathbf{G}_o)^{\frac{1}{2}} (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H \\ &= (\mathbf{G}_o^H \mathbf{G}_o)^{-\frac{1}{2}} \mathbf{G}_o^H \\ &= (\mathbf{V}_{\mathbf{G}_o} \Sigma_{\mathbf{G}_o}^{-2} \mathbf{V}_{\mathbf{G}_o}^H)^{-\frac{1}{2}} \mathbf{V}_{\mathbf{G}_o} \Sigma_{\mathbf{G}_o} \mathbf{U}_{\mathbf{G}_o}^H \\ &= \mathbf{V}_{\mathbf{G}_o} \Sigma_{\mathbf{G}_o}^{-1} \mathbf{V}_{\mathbf{G}_o}^H \mathbf{V}_{\mathbf{G}_o} \Sigma_{\mathbf{G}_o} \mathbf{U}_{\mathbf{G}_o}^H \\ &= \mathbf{V}_{\mathbf{G}_o} \mathbf{I} \mathbf{U}_{\mathbf{G}_o}^H. \end{aligned} \tag{4.5}$$

Thus, the singular values of \mathbf{W} are all 1, and the whitened transform coefficients are stable. The above analysis applies to any of the GLRT detectors, regardless of the underlying transform matrix \mathbf{R} : all the singular values of the “whitened transform matrix” $\mathbf{W}_{\mathbf{R}}$ are 1. Thus, the process of whitening the transform coefficients causes all of the GLRT statistics to be equally stable. With stability issues no longer a factor, the favorable properties of the OS-LSM-GT should surface (and they do, as was demonstrated in the last chapter).

Outside the realm of detection problems, however, there are other desirable qualities in a linear TFR. Some of these qualities include: low reconstruction error, accurate representation of the time-frequency content of signal components, and good localization (in the time-frequency plane) of signal components. The “transform-whitened” OS-LSM-GT, while it has good coefficient stability and good accuracy, has only fair localization and high reconstruction error. The next section investigates transform whitening and several other modifications to the OS-LSM-GT in terms of their ability to improve the coefficient stability of the OS-LSM-GT while also possessing the qualities listed above.

4.3 Methods of Improving Coefficient Stability in the OS-LSM-GT

4.3.1 Rank reduction

Rank reduction is an SVD-based method for improving the calculational stability of any operation involving matrix inversion. As the pseudoinverse in the OS-LSM-GT contains a matrix inverse, rank reduction can improve its coefficient stability. Consider the representation of the OS-LSM-GT of an observation \vec{y} , using the SVD of \mathbf{G}_o given in (4.3):

$$\begin{aligned}\vec{z} &= \mathbf{B}_o \vec{y} \\ &= \mathbf{V}_{\mathbf{G}_o} \Sigma_{\mathbf{G}_o}^{-1} \mathbf{U}_{\mathbf{G}_o}^H \vec{y}.\end{aligned}\quad (4.6)$$

For convenience, suppose that the MN singular values σ_i along the diagonal of $\Sigma_{\mathbf{G}_o}$ are ordered from largest to smallest, so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{MN}$. The corresponding MN singular values s_i along the diagonal of $\Sigma_{\mathbf{G}_o}^{-1}$ are then ordered in reverse as $s_1 \leq s_2 \leq \dots \leq s_{MN}$, since $s_i = \frac{1}{\sigma_i}$.

Given the above notation and given $r < MN$, the rank- r (*reduced-rank*) approximation of \vec{z} is given by

$$\vec{z} = \mathbf{V}_{\mathbf{G}_o} (\Sigma_{\mathbf{G}_o}^{-1})_r \mathbf{U}_{\mathbf{G}_o}^H \vec{y}, \quad (4.7)$$

where $(\Sigma_{\mathbf{G}_o}^{-1})_r$ is obtained from $\Sigma_{\mathbf{G}_o}^{-1}$ by setting $s_{r+1} = s_{r+2} = \dots = s_{MN} = 0$. It is well-known [TKK82, TS89] that $(\Sigma_{\mathbf{G}_o}^{-1})_r$ is the *best* rank- r approximation of $\Sigma_{\mathbf{G}_o}^{-1}$, in the sense that

$$\|\Sigma_{\mathbf{G}_o}^{-1} - (\Sigma_{\mathbf{G}_o}^{-1})_r\| \leq \|\Sigma_{\mathbf{G}_o}^{-1} - \mathbf{A}\| \quad (4.8)$$

for all rank- r matrices \mathbf{A} . The norm in (4.8) is the Frobenius norm described in the previous section.

The process of rank reduction increases the coefficient stability of the OS-LSM-GT by discarding the small singular values in $\Sigma_{\mathbf{G}_o}$ (or, equivalently, the large singular values in $\Sigma_{\mathbf{G}_o}^{-1}$) that cause the instability. Of course, there is a tradeoff involved in the process of rank reduction: as the rank r decreases, the stability of the OS-LSM-GT increases, but the mean-squared error $\|\vec{z} - \vec{y}\|$ (which, by the definition of the LSM, is *minimized* when there is no rank reduction) increases. Scharf [Sch91, TS89] calls this a *bias-variance* tradeoff.

The bias-variance tradeoff makes it important to choose the rank r of the reduced-rank solution carefully. Unfortunately, although some systematic methods of choosing r have been proposed [Sch91, section 9.16], the choice of r is essentially *ad hoc*.

Consider now Figure 4.3, which shows the rank-reduced OS-LSM-GT of the sample signals in the previous section. For convenience, the signal with (t, f) coordinates $(1, 1)$ will be referred to from now on as Signal 1, while the signal with (t, f) coordinates $(1.05, 1)$ will be referred to as Signal 2. The rank was chosen to be $r = 42$ (recall that $MN = 64$) for the plots of Figure 4.3. On the left-hand side of Figure 4.3, note that the rank-reduced transform coefficients of Signal 1 (no location mismatch) are a blurred version of the original OS-LSM-GT coefficients, but the one signal component is clearly depicted and centered at the correct (t, f) coordinates. Of course, rank reduction is not actually necessary for Signal 1, as the (unmodified) OS-LSM-GT coefficients already represent the desired solution.

On the right-hand side of Figure 4.3, note that the rank-reduced transform coefficients of Signal 2 (small arrival-time mismatch) are of the same order of magnitude as the actual signal component magnitude (which is 1). However, although the transform coefficients peak at the correct arrival-time ($t = 1$), they do not peak at the correct frequency ($f = 1$); instead there seem to be three components in the reduced-rank solution, none of them centered at $(1, 1)$. Changing the value of r changes the relative magnitudes of the reduced-rank transform coefficients, but no value of r ever produces one peak at the correct (t, f) coordinates. Thus, while greater stability is attained in the process of rank reduction, the sacrifice in accuracy is too great, at least for this simple example.

4.3.2 Regularization

Regularization is a method that attempts to improve the coefficient stability of the OS-LSM-GT by prescribing a penalty for high-frequency energy while finding the least-squared-error transform coefficients. This penalty takes the form of *constraints* on the signal $\vec{x} = \mathbf{G}_o \vec{c}$, so regularization is a problem in *constrained least squares* (CLS) [LHW94, Sch91].

To see how the constraints are imposed, recall that for the observation $\vec{y} = \mathbf{G}_o \vec{c} + \vec{n}$ within the PF framework, the LSM finds the solution \hat{c} that minimizes the squared-error

$$\|\vec{y} - \mathbf{G}_o \vec{c}\|,$$

where the norm is the 2-norm for vectors. Regularization seeks to find a “compromise” solution between the minimum squared error solution and a maximally smoothed solution satisfying

$$\mathbf{H} \mathbf{G}_o \vec{c} = \vec{0}, \quad (4.9)$$

where \mathbf{H} performs highpass filtering on the signal $\vec{x} = \mathbf{G}_o \vec{c}$.

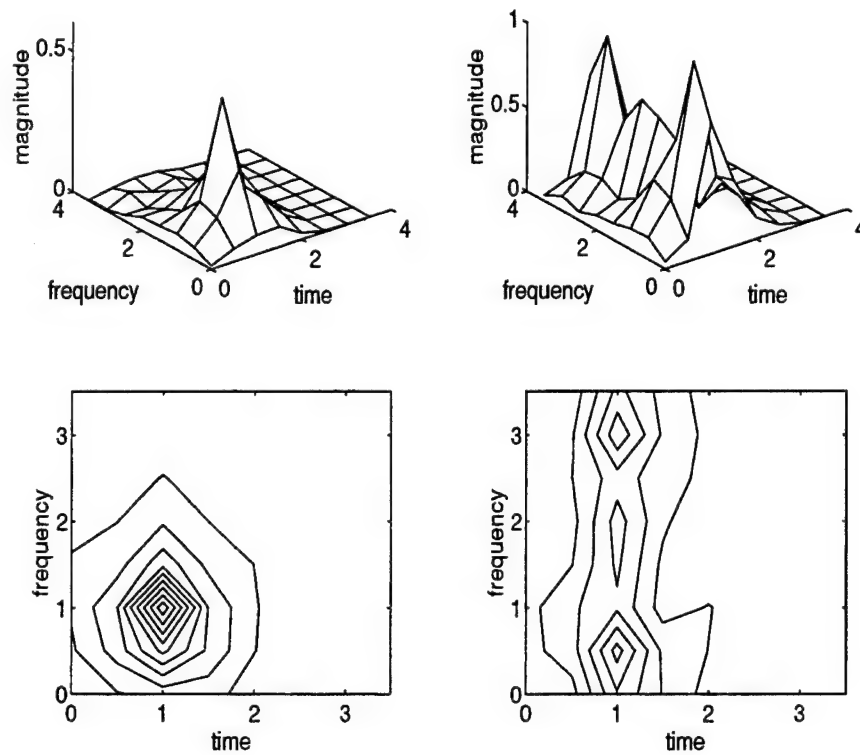


Figure 4.3: Left-hand side: magnitudes and contour plot of reduced-rank OS-LSM-GT coefficients of one-component signal, located at (1, 1). Right-hand side: magnitudes and contour plot of reduced-rank OS-LSM-GT coefficients of one-component signal, located at (1.05, 1). The rank reduction is performed at $r = 42$.

The highpass filtering operation is implemented by setting the elements of \mathbf{H} equal to the samples of a discrete-time highpass filter $h[k]$ as follows:

$$\mathbf{H}_{kl} = \begin{cases} h[k-l], & k-l \geq 0 \\ 0, & \text{else} \end{cases} \quad (4.10)$$

Thus, the k th element of the vector $\mathbf{H}\vec{x}$ is given by

$$\sum_{l=0}^{L-1} h[k-l]x[l],$$

which is the familiar discrete-time convolution sum.

The CLS problem arising from regularization, then, is

$$\min_{\vec{c}} \|\vec{y} - \mathbf{G}_o \vec{c}\| + \lambda_{reg} \|\mathbf{H} \mathbf{G}_o \vec{c}\|. \quad (4.11)$$

The solution to (4.11) is given by

$$\hat{c} = \vec{z} = (\mathbf{G}_o^H \mathbf{G}_o + \lambda_{reg} \mathbf{G}_o^H \mathbf{H}^H \mathbf{H} \mathbf{G}_o)^{-1} \mathbf{G}_o \vec{y}, \quad (4.12)$$

where the *regularization parameter* λ_{reg} controls the degree of solution-smoothing. When $\lambda_{reg} \rightarrow 0$, there is no smoothing and the CLS solution is the same as the unconstrained LSM solution. When $\lambda_{reg} \rightarrow \infty$, smoothing is dominant but the CLS solution loses accuracy.

Figure 4.4 below illustrates the use of regularization on Signal 2. Application of regularization to Signal 1 is unnecessary, as the unconstrained LSM solution ($\lambda_{reg} = 0$) gives the desired result. The highpass filter $h[k]$ used to generate \mathbf{H} is an 8th-order, recursive IIR filter with stopband frequency $f_s = 4$ Hz. The choice of $f_s = 4$ Hz makes use of the PF framework assumption that the signal does not have any frequency components above 3.5 Hz in the oversampled case.

In the top part of Figure 4.4, the transform coefficients are generated using $\lambda_{reg} = 10$, signifying a minimal amount of smoothing. Consequently, the solution resembles the unconstrained LSM solution, but the transform coefficient magnitudes are one order of magnitude smaller. Even with this minimal amount of smoothing, the transform localizes the one signal component, although to the wrong (t, f) coordinates $((1, 2)$ instead of $(1, 1)$). In the middle part of Figure 4.4, $\lambda_{reg} = 10^4$. With this amount of smoothing, the transform coefficient magnitudes are only one order of magnitude larger than the actual signal component magnitude of 1; moreover, the transform localizes the one signal component to $(1, 1.5)$, which is adjacent to the ideal location of $(1, 1)$ (recall that the actual (t, f) coordinates of the signal component in Signal 2 are $(1.05, 1)$). In the bottom part of Figure 4.4, $\lambda_{reg} = 10^5$, which produces transform coefficient magnitudes on the same order of magnitude as the actual signal component magnitude. However, the corresponding contour plot shows

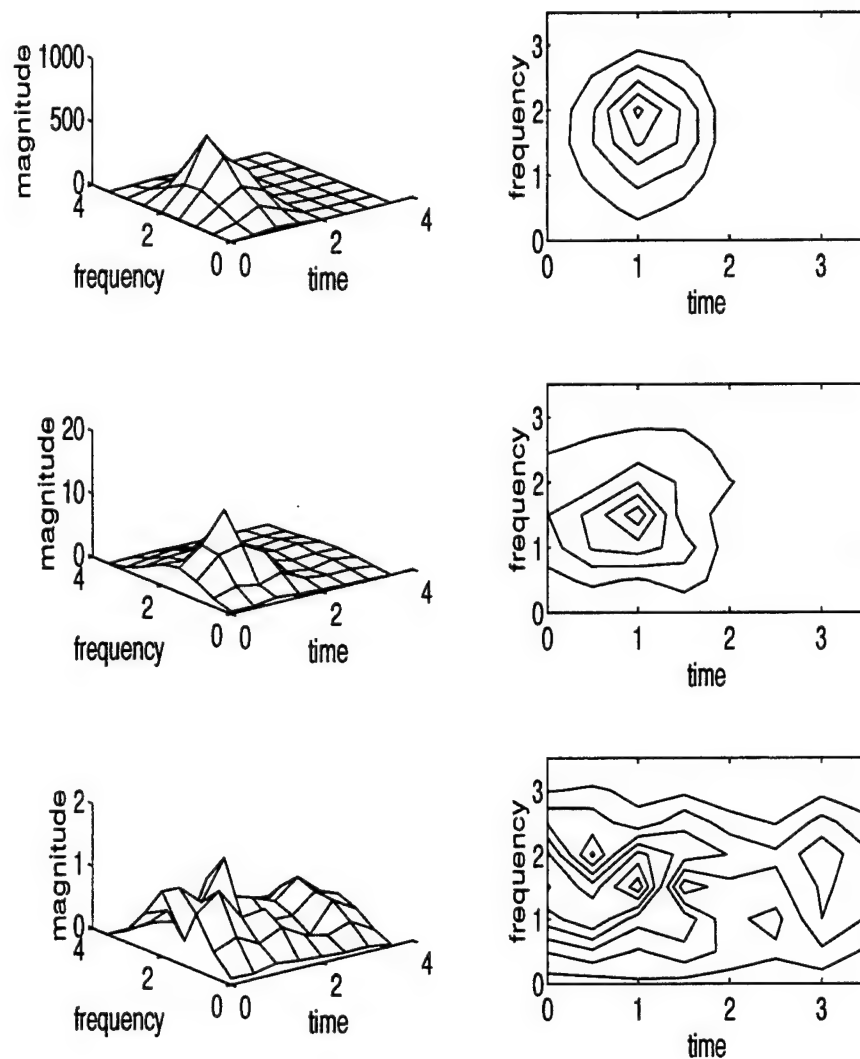


Figure 4.4: Regularized OS-LSM-GT coefficients of one-component signal, located at (1.05, 1). Top: coefficient magnitudes and contour plot for $\lambda_{reg} = 10$. Middle: coefficient magnitudes and contour plot for $\lambda_{reg} = 10^4$. Bottom: coefficient magnitudes and contour plot for $\lambda_{reg} = 10^5$.

that the transform coefficients no longer localize the signal component accurately; in fact, they do not even localize to one signal component.

In the process of regularization, then, there is a tradeoff – controlled by the value of λ_{reg} – between smoothing and accuracy. However, it is clear from the plots of Figure 4.4 that no value of λ_{reg} produces transform coefficient magnitudes on the same order of magnitude as the actual signal component *and* localizes the signal component to the correct (t, f) coordinates. Moreover, in many problems one does not know what the “correct” solution should be, so the proper selection of λ_{reg} is something of an *ad hoc* process, just as the selection of the rank r in rank reduction is somewhat *ad hoc*.

4.3.3 Principal components method

The **principal components method** (PCM), first proposed by Tufts, Kumaresan, and KIRSTEINS [TKK82, TK82], also uses rank reduction, but as a preprocessing tool rather than as a direct means of improving the coefficient stability of the OS-LSM-GT. The PCM prescribes the use of the observation vector $\vec{y} = \mathbf{G}_o \vec{c} + \vec{v}$ to form a *forward-backward linear prediction* (FBLP) data matrix. For a length L observation vector \vec{y} , the FBLP matrix \mathbf{Y} is given by

$$\mathbf{Y} = \begin{bmatrix} y[n] & y[n-1] & \dots & y[1] \\ y[n+1] & y[n] & \dots & y[2] \\ \vdots & \vdots & \vdots & \vdots \\ y[L] & y[L-1] & \dots & y[L-n+1] \\ y^H[1] & y^H[2] & \dots & y^H[n] \\ y^H[2] & y^H[3] & \dots & y^H[n+1] \\ \vdots & \vdots & \vdots & \vdots \\ y^H[L-n+1] & y^H[L-n+2] & \dots & y^H[L] \end{bmatrix}, \quad (4.13)$$

where n is the assumed length of an adaptive linear prediction filter. The dimensions of \mathbf{Y} are $([2L - 2n + 2] \times n)$.

Now consider the SVD of \mathbf{Y} :

$$\mathbf{Y} = \mathbf{U}_Y \Sigma_Y \mathbf{V}_Y^H. \quad (4.14)$$

The principal components method approximates \mathbf{Y} using the reduced-rank matrix $(\Sigma_Y)_r$:

$$\mathbf{Y}_r = \mathbf{U}_Y (\Sigma_Y)_r \mathbf{V}_Y^H. \quad (4.15)$$




Figure 4.5: Comparison of one-component signal located at (1.05, 1) and its rank-1 approximation using PCM (filter length 8). Only the real parts of the time histories are shown.

The r singular values that are not discarded correspond to the *principal components* of \mathbf{Y} (or \vec{y}). Since there is a one-to-one correspondence between a data vector and its FBLP data matrix, the approximate observation \vec{y}_r can then be reconstructed from \mathbf{Y}_r . Thus, the OS-LSM-GT with PCM is given by

$$\vec{z} = \mathbf{B}_o \vec{y}_r. \quad (4.16)$$

The PCM improves the coefficient stability of the OS-LSM-GT indirectly, by “simplifying” the observation vector \vec{y} so that the transform is performed on the most meaningful portion of the observation data. If \vec{y} is derived from a signal with k components, it is natural to approximate \vec{y} using PCM by \vec{y}_r , with $r = k$.

Figures 4.5 and 4.6 illustrate the application of PCM to the OS-LSM-GT of Signal 1 and Signal 2. The filter length in (4.13) was taken to be 8, and the observation vector \vec{y} was approximated using just the first principal component. Figure 4.5 depicts the real parts of the time histories of Signal 2 and of the rank-1 approximation to Signal 2. Because the largest singular value in $\Sigma_{\mathbf{Y}}$ is much larger than the other seven singular values, there is little difference in approximating \vec{y} by its first principal component. The major difference evident in the approximation is that the jump discontinuity of the transient signal component appears “smoother.”

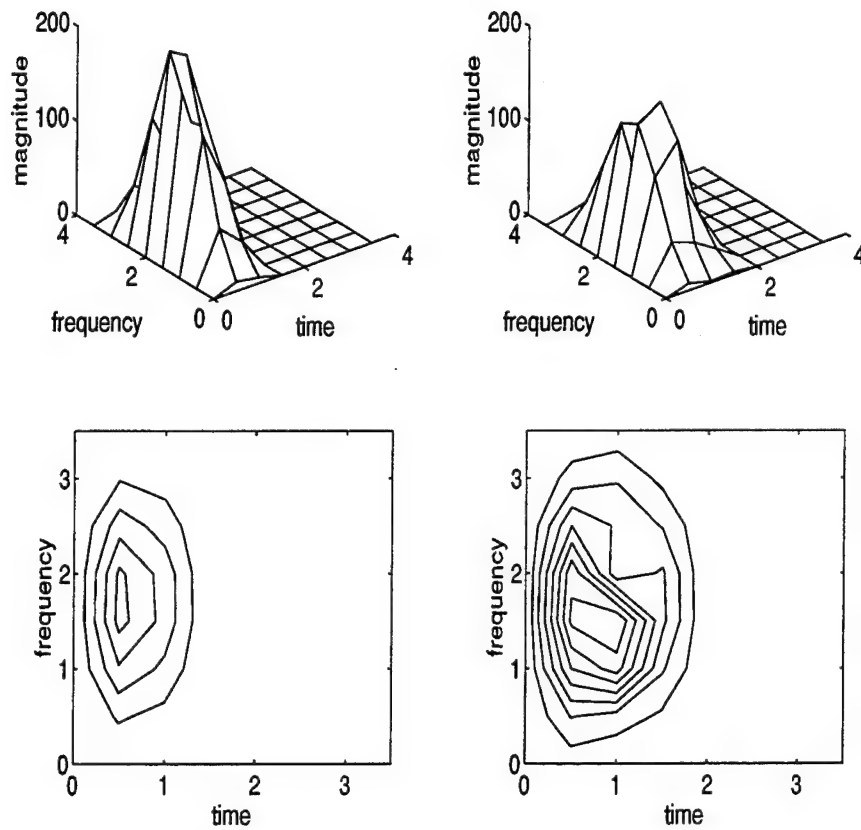


Figure 4.6: Left-hand side: magnitudes and contour plot of OS-LSM-GT coefficients for rank-1 approximation of signal located at (1, 1). Right-hand side: magnitudes and contour plot of OS-LSM-GT coefficients for rank-1 approximation of signal located at (1.05, 1).

Figure 4.6 shows the coefficient magnitudes and contour plots of the OS-LSM-GT, performed on the rank-1 approximations of Signal 1 (left-hand side) and Signal 2 (right-hand side). Recall that the OS-LSM-GT applied to Signal 1 without PCM produces a completely accurate and localized solution; however, the OS-LSM-GT applied to the rank-1 approximation of Signal 1, while fairly localized, is no longer accurate. The OS-LSM-GT applied to the rank-1 approximation of Signal 2 produces coefficient magnitudes one order of magnitude closer to that of the actual signal component than the coefficient magnitudes produced by applying the OS-LSM-GT to Signal 2 directly; however, the localization is again inaccurate. Thus, as with rank reduction and with regularization, the principal component fails to produce an accurate and localized modification of the OS-LSM-GT of Signal 2.

4.3.4 LSM solution with energy constraints

While regularization solves a CLS problem with a linear equality constraint, the **energy-constrained LSM** (EC-LSM) solves a CLS problem with a quadratic inequality constraint:

$$\begin{aligned} \min_{\vec{c}} \quad & \|\vec{y} - \mathbf{G}_o \vec{c}\|, \\ \text{subject to} \quad & \|\vec{c}\| \leq E_0, \end{aligned} \tag{4.17}$$

where E_0 is the maximum allowable energy of the transform coefficient vector \vec{c} . Thus, the EC-LSM attempts to find the least squares estimate of \vec{c} among all vectors that satisfy the energy constraint. At the very least, this prevents the calculation of transform coefficients whose magnitudes are several orders of magnitude above the actual signal component magnitudes (as happens with the unmodified OS-LSM-GT). However, the EC-LSM does require fairly accurate *a priori* knowledge of E_0 .

The solution of (4.17) is given in [GL83, section 12.1] and involves Lagrange multipliers and SVD; the algorithm will not be repeated here. Figure 4.7 illustrates the application of the EC-LSM to Signal 2, using $E_0 = 2$ (the actual energy of the coefficient vector for Signal 2 is 1)³. In Figure 4.7, note that the magnitudes of the transform coefficients are of the same order of magnitude as that of the actual signal component; also note that the transform coefficients are localized to one main component located *near* the ideal (t, f) coordinates $((1, 0.5)$ instead of $(1, 1)$) and to several smaller components with correct arrival times ($= 1$) but incorrect frequencies. Thus, while the EC-LSM provides a better solution than any of the methods mentioned thus far, with good localization and accurate coefficient magnitudes, it

³Application of the EC-LSM to Signal 1 produces a solution identical to the unconstrained OS-LSM-GT coefficients, because the latter solution produces 0 squared-error while satisfying the energy constraint.




Figure 4.7: Magnitudes and contour plot of EC-LSM coefficients for signal located at (1.05, 1).

still does not provide a completely accurate location for the signal component in Signal 2.

4.3.5 LSM solution with transform “whitening”

The final method of improving the OS-LSM-GT that will be examined in this section is the use of the “whitened” transform vector $\vec{\zeta}$, defined previously as

$$\vec{\zeta} = (\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}} \vec{z} \quad (4.18)$$

$$= (\mathbf{G}_o^H \mathbf{G}_o)^{\frac{1}{2}} \vec{z}, \quad (4.19)$$

in place of the “unwhitened” transform vector \vec{z} .

As noted in section 3.7, the covariance matrix of $\vec{\zeta}$ is $2\mathbf{I}$ (hence the term, “whitened” transform vector). And as derived in section 4.2, the singular values of $\vec{\zeta}$ are all 1. Thus, whitening the transform greatly increases the coefficient stability of the OS-LSM-GT.

However, there is a blurring effect that accompanies multiplication by the whitening matrix $(\mathbf{B}_o \mathbf{B}_o^H)^{-\frac{1}{2}}$. This blurring is nonexistent if the transform basis functions are orthogonal (the Gabor logons are highly nonorthogonal, especially in the oversampled case; this is true regardless of the window function, as the Gabor expansion always contains several logons that have the same arrival time but different frequencies).

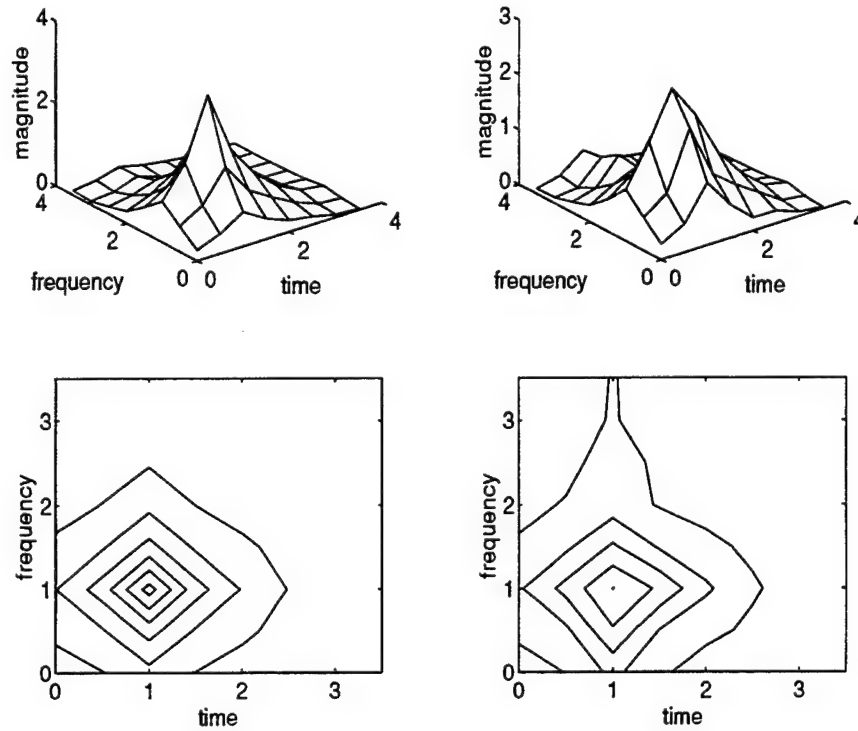


Figure 4.8: Left-hand side: magnitudes and contour plot of whitened OS-LSM-GT coefficients for signal located at $(1, 1)$. Right-hand side: magnitudes and contour plot of whitened OS-LSM-GT coefficients for signal located at $(1.05, 1)$.

Figure 4.8 illustrates the effectiveness of the whitened OS-LSM-GT. The left-hand side of Figure 4.8 shows the magnitudes and corresponding contour plot of the whitened OS-LSM-GT coefficients for Signal 1. Note that while these plots are a blurred version of the corresponding unwhitened OS-LSM-GT coefficients, the signal component is still fairly well-localized and still represented accurately. The right-hand side of Figure 4.8 shows the magnitudes and corresponding contour plot of the whitened OS-LSM-GT coefficients for Signal 2. Unlike any of the methods mentioned above, transform whitening localizes the one signal component to the ideal (t, f) coordinates $(1, 1)$, even though there is some blurring.

The whitened OS-LSM-GT can be combined with the PCM for even more effective results. Figure 4.9 shows the magnitudes and corresponding contour plots of the whitened OS-LSM-GT coefficients for the rank-1 approximations of Signals 1 and 2. Examination of the four plots reveals that the whitened OS-LSM-GT, when used

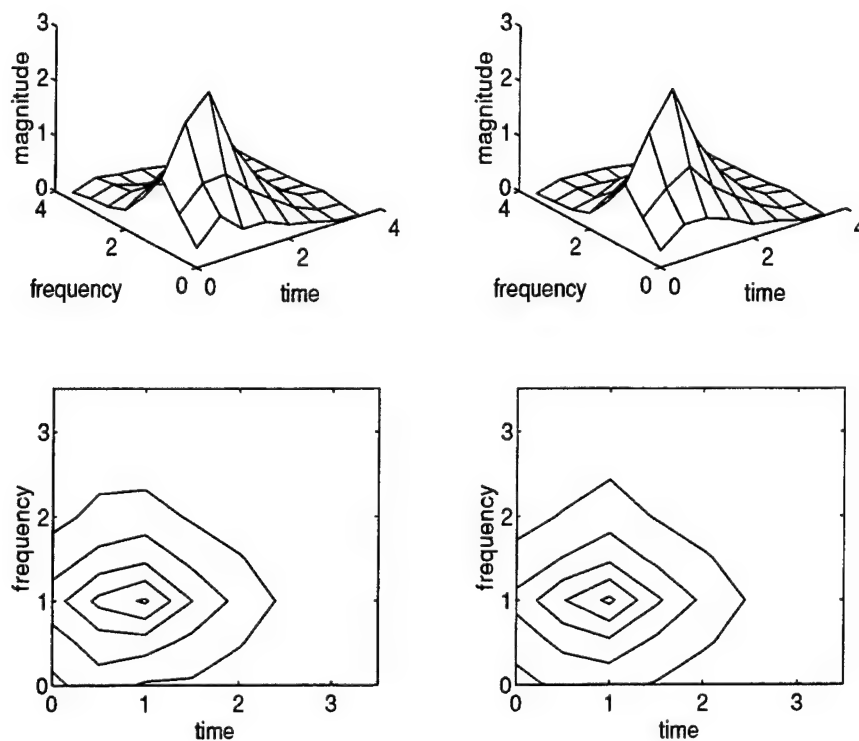


Figure 4.9: Left-hand side: magnitudes and contour plot of whitened OS-LSM-GT coefficients for rank-1 approximation (using PCM) of signal located at (1, 1). Right-hand side: magnitudes and contour plot of whitened OS-LSM-GT coefficients for rank-1 approximation (using PCM) of signal located at (1.05, 1).

with the PCM, shows very little sensitivity to the shift in component arrival time between Signals 1 and 2.

Note that the whitened OS-LSM-GT is somewhat similar to the oversampled SSTFT with either rectangular window or one-sided exponential window (see Figures 3.9 and 3.10), in that all three of the transforms provide fairly accurate but not very localized representations of the signal component.

4.4 Performance Analysis of Modifications to the OS-LSM-GT

This section quantifies the abilities of the five modifications to the OS-LSM-GT mentioned above to represent accurately some simple signals. Three performance measures will be used, each of which measures different aspects of the accuracies of the respective modifications. The analysis will be performed on three signals: Signal 1, Signal 2, and a combination of Signal 2 with noise.

4.4.1 Performance measures

The three performance measures that will be used in the analysis are as follows:

1. the *normalized reconstruction error* (NRE), derived from the reconstructed signal $\hat{y} = G_o \hat{c} = G_o \tilde{z}$:

$$NRE = \frac{\|G_o \tilde{z} - \tilde{y}\|}{\|\tilde{y}\|}. \quad (4.20)$$

NRE measures how accurately the transform coefficients \tilde{z} can reconstruct the observation \tilde{y} . $NRE > 1$ generally indicates poor reconstruction accuracy (the norm of the error is greater than the norm of the original observation if this is the case).

2. the *concentration* (CONC) of the transform coefficients:

$$CONC = \frac{\tilde{z}_{sig}^H \tilde{z}_{sig}}{\tilde{z}^H \tilde{z}}, \quad (4.21)$$

where \tilde{z}_{sig} is comprised of any transform coefficients whose (t, f) coordinates are at or adjacent to the (t, f) coordinates of the actual signal components. For Signal 1, for example, \tilde{z}_{sig} is comprised of the transform coefficients with coordinates (0.5, 0.5), (0.5, 1), (0.5, 1.5), (1, 0.5), (1, 1), (1, 1.5), (1.5, 0.5), (1.5, 1), and (1.5, 1.5). CONC measures how well the transform coefficients are localized to the actual signal component locations by taking the ratio of the energy of the transform coefficients at or near the supposed (t, f) coordinates of the actual signal components to the energy of *all* the transform coefficients. $CONC = 1$ indicates perfect localization; $CONC \approx 0$ indicates poor localization (i.e., the component is localized inaccurately, or the component is not localized).

3. the correlation coefficient (CC) between the transform coefficients and the actual signal component locations. This is defined in [GDL94] in terms of the *actual* transform coefficient magnitudes $d_{mn} = |c_{mn}|$ and the *estimated* transform coefficient magnitudes $\hat{d}_{mn} = |\hat{c}_{mn}|$ (recall that in the Gabor expansion (2.31), m corresponds to frequency and n corresponds to time):

$$CC = \frac{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} d_{mn} \hat{d}_{mn} - \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} d_{mn} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \hat{d}_{mn}}{\sqrt{\left\{ \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} d_{mn}^2 - \frac{1}{MN} \left[\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} d_{mn} \right]^2 \right\} \left\{ \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \hat{d}_{mn}^2 - \frac{1}{MN} \left[\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \hat{d}_{mn} \right]^2 \right\}}} \quad (4.22)$$

It is necessary to use the coefficient magnitudes in (4.22) instead of the complex-valued coefficients themselves because CC is defined for real-valued coefficients. CC measures the similarity between \vec{d}_{mn} and $\hat{\vec{d}}_{mn}$. It is well known that $-1 \leq CC \leq 1$; $CC = 1$ indicates that \vec{d} and $\hat{\vec{d}}$ are perfectly correlated, while $CC \approx 0$ indicates that \vec{d} and $\hat{\vec{d}}$ are uncorrelated.

4.4.2 Analysis results

Table 4.2 below lists the values of the three performance measures attained by the OS-LSM-GT and by several of its modifications when applied to Signal 1. Recall that the samples of Signal 1 are given by

$$x[k] = x(k\Delta) = \sqrt{2}e^{-(k\Delta-1)}e^{j2\pi k\Delta}u(k\Delta-1), \quad k = 0, 1, \dots, L-1, \quad (4.23)$$

with $\Delta = \frac{1}{32}$. Signal 1 is thus a one-component signal with one-sided exponential window (parameter $\lambda = 1$) located at (t, f) coordinates $(1, 1)$. The transforms and modifications included in Table 4.2 are as follows:

1. the OS-LSM-GT (without modification)
2. the OS-LSM-GT with rank reduction (symbol: RR), using rank $r = 42$
3. the OS-LSM-GT with regularization (symbol: REG), using the highpass filter described in subsection 4.3.2 above and regularization parameter $\lambda_{reg} = 10^4$
4. the OS-LSM-GT with the principal components method (symbol: PCM), using a rank-1 approximation (i.e., only the first principal component) of the observation \vec{y}

5. the OS-LSM-GT with energy constraints (symbol: EC), using an energy constraint of $E_0 = 2$ on the energy of the coefficient vector \vec{c}
6. the OS-LSM-GT with transform whitening (symbol: TW)
7. the OS-LSM-GT with transform whitening *and* with the principal components method (symbol: TW-PCM), using a rank-1 approximation of \vec{y}
8. the oversampled SSTFT with 1-second rectangular window (symbol: OS-STFT-R1)
9. the oversampled SSTFT with 0.5-second rectangular window (symbol: OS-STFT-R2)
10. the oversampled SSTFT with one-sided exponential window having parameter $\lambda = 1$ (symbol: OS-STFT-EW).

Note that three forms of the oversampled SSTFT are included for comparison with the OS-LSM-GT and its various modifications.

Transform	Performance measure		
	NRE	CONC	CC
1.OS-LSM-GT	ϵ	1	1
2.RR	0.0015	0.8846	0.7436
3.REG	0.7571	0.5923	0.2426
4.PCM	0.3540	0.5308	0.1703
5.EC	ϵ	1	1
6.TW	10.2770	0.7701	0.6199
7.TW-PCM	9.4307	0.7299	0.5256
8.OS-STFT-R1	2.6018	0.7864	0.4994
9.OS-STFT-R2	1.2440	0.6334	0.4326
10.OS-STFT-EW	3.1407	0.7124	0.5881

Table 4.2: Comparison of performance measures for various transforms of Signal 1. This is the no mismatch case. Any value of NRE less than 10^{-9} is denoted by ϵ .

The values of the performance measures in Table 4.2 confirm many of the observations of the previous section. Note that Signal 1 does not contain any information mismatch, so the OS-LSM-GT performs perfectly well for this case. A few observations can be made about the performances of the other transforms and modifications:

- The energy-constrained (EC) coefficients are the same as the (unmodified) OS-LSM-GT coefficients for no-mismatch signals such as Signal 1.
- RR performs quite well, since it modifies the already perfect OS-LSM-GT coefficients.
- REG and PCM both have low CC values: REG, because there is too much smoothing; and PCM, because the rank-1 approximation *creates* an information mismatch although there is none in the original signal.
- TW and TW-PCM localize the component accurately (note the high CONC and CC values) but are *extremely* inaccurate in reconstructing the signal (very high values of NRE). This is not surprising, as the process of transform whitening increases coefficient stability at the expense of reconstruction accuracy.
- The SSTFTs all perform about the same, with relatively accurate localization but poor reconstruction accuracy.

Table 4.3 below lists the values of the three performance measures attained by the same ten transforms as in Table 4.2 when applied to Signal 2. The time samples of Signal 2 are given by

$$x[k] = x(k\Delta) = \sqrt{2}e^{-(k\Delta-1.05)}e^{j2\pi(k\Delta-1.05)}u(k\Delta-1.05), \quad k = 0, 1, \dots, L-1, \quad (4.24)$$

with $\Delta = \frac{1}{32}$. Signal 2 is thus a one-component signal with one-sided exponential window (parameter $\lambda = 1$) located at (t, f) coordinates $(1.05, 1)$. Because Signal 2 represents a very slight modification of Signal 1, the actual transform coefficient magnitudes d_{mn} are assumed to be the same as those for Signal 1:

$$d_{mn} = \begin{cases} 1, & m = n = 2 \\ 0, & \text{else} \end{cases} \quad (4.25)$$

Recall that $m = n = 2$ in the Gabor expansion with $\alpha = \beta = 0.5$ corresponds to (t, f) coordinates of $(1, 1)$.

The values of the performance measures in Table 4.3 again confirm many of the observations of the previous section:

- The OS-LSM-GT still produces a small reconstruction error (which it will always do by definition), but it no longer localizes the signal component perfectly or even satisfactorily. Figure 4.2 from section 4.2 confirms this observation.

Transform	Performance measure		
	NRE	CONC	CC
1.OS-LSM-GT	0.0951	0.5027	0.2939
2.RR	0.1269	0.3949	0.1971
3.REG	0.7564	0.6745	0.3223
4.PCM	0.3560	0.6867	0.3605
5.EC	0.1264	0.5623	0.1125
6.TW	10.0222	0.7029	0.4885
7.TW-PCM	9.5981	0.7389	0.5445
8.OS-STFT-R1	2.5792	0.7657	0.4836
9.OS-STFT-R2	1.2563	0.5908	0.3885
10.OS-STFT-EW	3.0490	0.7000	0.5682

Table 4.3: Comparison of performance measures for various transforms of Signal 2. There is a small arrival time mismatch.

- None of the modifications to the OS-LSM-GT suggested in section 4.3 localize the signal component accurately, with all of them producing CC values below 0.5 except TW-PCM.
- The values of the performance measures for the SSTFTs of Signal 2 all remain about the same as the values for Signal 1. This is not surprising, as the SSTFT is an inner-product-based transform whose values should not change very much with the slight modification from Signal 1 to Signal 2.

Table 4.4 below lists the values of the three performance measures attained by the same ten transforms when applied to Signal 2 with complex Gaussian noise ($\sim \mathcal{N}(\vec{0}, \mathbf{I})$) added in. The actual transform coefficient magnitudes d_{mn} are still assumed to be the same as those for Signal 1 (see (4.25)).

It is clear from the first several rows of Table 4.4 that neither the OS-LSM-GT nor most of the modifications suggested in section 4.3 are able to localize the signal component with any great degree of accuracy. The SSTFTs and the TW-PCM provide fairly accurate localization of the signal component, but – as is the case with all three examples – they do not provide accurate reconstruction of the signal from the transform coefficients.

The following general conclusions can be drawn from the results of this section and the observations of the previous section for oversampled linear TFRs within the PF framework:

Transform	Performance measure		
	NRE	CONC	CC
1.OS-LSM-GT	0.7628	0.3770	0.2650
2.RR	0.8191	0.1963	0.0427
3.REG	0.9492	0.3373	0.1811
4.PCM	0.9254	0.3268	0.1502
5.EC	0.8142	0.1668	0.0239
6.TW	4.9279	0.4552	0.3789
7.TW-PCM	5.0365	0.5423	0.5344
8.OS-STFT-R1	1.4822	0.6260	0.4772
9.OS-STFT-R2	1.0497	0.5391	0.4508
10.OS-STFT-EW	1.6171	0.5900	0.5717

Table 4.4: Comparison of performance measures for various transforms of Signal 2 added to complex noise $\sim \mathcal{N}(\vec{0}, \mathbf{I})$. The SNR is approximately -9.0 dB.

- If the main goal of the oversampled transform is to reconstruct the signal as accurately as possible, then the OS-LSM-GT should be used. However, the OS-LSM-GT does not give an accurate indication of the actual time-frequency content of the signal.
- If the main goal of the oversampled transform is to represent the actual time-frequency content of the signal accurately and with fairly good localization, then either the OS-LSM-GT with transform whitening and PCM or a version of the SSTFT should be used. However, neither the TW-PCM nor any of the SSTFTs can accurately reconstruct the signal.
- Rank reduction, regularization, and energy-constrained solutions are very ineffective in providing accurate indications of the actual time-frequency content of noisy or information-mismatched signals, although all of these methods do increase the coefficient stability of the OS-LSM-GT. However, this increase in stability does not lead to a corresponding increase in transform accuracy. Of all the methods suggested in section 4.3 to increase the coefficient stability of the OS-LSM-GT, only transform whitening is able to localize signal components with any degree of accuracy.

4.5 Wexler-Raz (WR) Framework for Discrete-Time Gabor Transforms

Up to this point, the transforms and modifications examined in this chapter have been designed to implement the oversampled Gabor transform within the PF framework. One of the important assumptions of the PF framework is assumption A1, that the number of possible nonzero Gabor transform coefficients MN is less than the length L of the signal \tilde{x} . This section and the next two investigate a framework originally proposed by Wexler and Raz [WR90] in which the oversampled Gabor transform is calculated under the assumption that $MN > L$.

The Wexler-Raz (WR) framework represents the discrete Gabor expansion of a length- L signal \tilde{x} as

$$x[k] = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g[k - N_1 n] e^{\frac{j2\pi m M_1 k}{L}}, \quad (4.26)$$

$$= \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g_{mn}[k], \quad k = 0, 1, \dots, L-1, \quad (4.27)$$

where M_1 and N_1 represent the intervals of frequency and time, respectively, at which the Gabor transform is calculated and where $g_{mn}[k] = g[k - N_1 n] e^{\frac{j2\pi m M_1 k}{L}}$. The sampling frequency $f_s = \frac{1}{\Delta}$ is implicit in (4.26). In the WR framework, it is assumed that

$$MM_1 = NN_1 = L. \quad (4.28)$$

Stable reconstruction of the signal \tilde{x} from its Gabor transform coefficients requires that $M_1 N_1 \leq L$ or that, equivalently,

$$MN \geq L. \quad (4.29)$$

The case where $MN = L$ is the critical sampling case in the WR framework, while the case where $MN > L$ is the oversampling case.

The Gabor expansion in (4.26) is similar to the discrete-time Gabor expansion given in (3.2), which is repeated here for reference:

$$x[k] = x(k\Delta) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} g(k\Delta - n\alpha) e^{j2\pi m \beta k \Delta}, \quad k = 0, 1, \dots, L-1, \quad (4.30)$$

where the trailing $n\alpha$ in the complex exponential term of (3.2) has been discarded

to allow a direct comparison between (4.26) and (4.30)⁴. By comparing (4.26) and (4.30), one can rewrite the assumptions of the WR framework given by (4.28) in terms of the more familiar parameters α and β :

$$N\alpha = L\Delta \quad (4.31)$$

$$M\beta = \frac{1}{\Delta} \quad (4.32)$$

$$MN\alpha\beta = L. \quad (4.33)$$

These three assumptions will be labeled *B1*, *B2*, and *B3*, respectively; assumption *B3* is derived by multiplying *B1* and *B2* together. Inspection of assumption *B3* (4.33) reveals that the stable reconstruction condition (4.29) is equivalent to the familiar $\alpha\beta \leq 1$. In fact, *B3* shows that oversampling in the WR framework, $MN > L$, occurs $\Leftrightarrow \alpha\beta < 1$. Thus, the two definitions of oversampling are equivalent, at least within the WR framework.

The assumption (4.28), while necessary in Wexler and Raz's derivation of the discrete-time Gabor transform [WR90, Appendix D], leads to unnecessary restrictions in certain situations; this is best illustrated by an example. Suppose that one is given the samples of an 8 second signal, sampled every $\Delta = \frac{1}{32}s$; thus, there are $L = 256$ samples in the signal. For this signal, one wishes to calculate the Gabor transform coefficients, oversampled with $\alpha = \beta = 0.5$. Furthermore, suppose that it is known *a priori* that the maximum arrival time and maximum modulation frequency are both 3.5 for any of the signal components. Within the PF framework, one would set $M = N = 8$ and proceed to find the solution to the matrix equation $\vec{x} = \mathbf{G}_o \vec{c}$, with \mathbf{G}_o of dimension (256×64) . $MN = 64$ Gabor transform coefficients are calculated. These are the exact parameters within the PF framework of the transient detection problem that was analyzed in detail in Chapter 3.

However, note that assumption *B1* (4.31) of the WR framework dictates that N be set to $N = \frac{L\Delta}{\alpha} = 16$. This prescribes the calculation of Gabor transform coefficients for logons with arrival times all the way up to $(N - 1)\alpha = 7.5$. The latter would be appropriate if one were given an 8 second signal *without* any *a priori* information about the possible component arrival times. However, the information is given that the maximum possible arrival time for any signal component is 3.5. From a practical standpoint, then, it is unnecessary to set $N = 16$, because it is known *a priori* that any transform coefficients with time index $n > 7$ are zero.

Similarly, note that assumption *B2* (4.32) of the WR framework dictates that M

⁴Alternatively, one could substitute $(k - N_1n)$ for k in the complex exponential term of (4.26). The term in question represents a constant phase, so it can be incorporated into the coefficients C_{mn} .

be set to $M = \frac{1}{\beta\Delta} = 64$. This prescribes the calculation of Gabor transform coefficients for logons with modulation frequencies all the way up to $(M-1)\beta = 31.5$. The latter would be appropriate if one were given a signal sampled at $f_s = 32\text{Hz}$ *without* any *a priori* information about the possible component modulation frequencies, as 31.5 Hz is the maximum frequency on the oversampled Gabor lattice ($\alpha = \beta = 0.5$) that can be represented unaliased in this case⁵. However, the information is given that the maximum possible modulation frequency for any signal component is 3.5. Again, it is unnecessary from a practical standpoint to set $M = 64$, because it is known *a priori* that any transform coefficients with frequency index $m > 7$ are zero.

Within the WR framework for the example above, then, one has to set $N = 16$ and $M = 64$ and then find the solution to the matrix equation

$$\vec{x} = \mathbf{G}_o \vec{c}, \quad (4.34)$$

with \mathbf{G}_o of dimension (256×1024) . $MN = 1024$ Gabor transform coefficients are calculated. The main difference in calculating the oversampled Gabor transform within the PF and WR frameworks lies essentially with the *number* of transform coefficients that need to be calculated. As the example above shows, it is extraneous to calculate the additional transform coefficients required by the WR framework-assumption (4.28) when *a priori* information is given about the signal. In other words, it is only necessary to calculate the number of transform coefficients dictated by the WR framework when one knows *nothing* about the signal components. Nevertheless, the WR framework will be explored further in the next two sections because a great deal of recent work [Bal92, BO92, Orr91, Orr92, Orr93a, QC93, QC92, QC94a, QCL92, WR90] has focused on the calculation of the Gabor transform within the WR framework.

4.6 Calculation of the Oversampled Gabor Transform in the WR Framework

In the matrix equation (4.34) for the Gabor expansion within the WR framework, it is important to note that the signal subspace matrix \mathbf{G}_o has more columns (MN) than rows (L) in the oversampled case. Indeed, when the oversampling parameters are $\alpha = \beta = 0.5$, assumption B3 (4.33) dictates that \mathbf{G}_o has four times as many columns as rows. Thus, the matrix equation (4.34) represents an *underdetermined* linear system. This is exactly the opposite of the PF framework, where it is always

⁵A sampling rate of $f_s = 32\text{ Hz}$ corresponds to a Nyquist rate of 32 Hz when only positive frequencies are allowed, as is the case here.

assumed that $MN < L$ (and usually $MN \ll L$), so that the corresponding matrix equation for the Gabor expansion represents an *overdetermined* linear system.

Because the linear system (4.34) is underdetermined for the oversampled Gabor transform within the WR framework, there are an infinite number of *exact* solutions⁶. Following is a discussion of several methods of selecting the “best” solution among the infinite set of exact solutions.

4.6.1 Periodized biorthogonal function method

In [WR90], Wexler and Raz proposed the use of a discretized and periodized biorthogonal function to calculate the Gabor coefficients C_{mn} :

$$C_{mn} = \sum_{k=0}^{L-1} x[k] \tilde{\gamma}_{mn}^*[k] \quad (4.35)$$

$$= \sum_{k=0}^{L-1} x[k] \tilde{\gamma}^*[k - nN_1] e^{\frac{j2\pi m M_1 k}{L}} \quad (4.36)$$

$$= \langle x, \tilde{\gamma}_{mn} \rangle, \quad (4.37)$$

where $\tilde{\gamma}_{mn}$ is periodic with period L . The periodicity of $\tilde{\gamma}_{mn}$ causes the transform coefficients C_{mn} to be doubly periodic, with period M in ‘m’ (frequency) and period N in ‘n’ (time).

The discrete-time biorthogonal function γ must satisfy the discrete-time biorthogonality condition, which was determined by Wexler and Raz [WR90] to be

$$\sum_{k=0}^{L-1} \tilde{g}[k + nM] e^{-\frac{j2\pi m N k}{L}} \tilde{\gamma}^*[k] = \frac{L}{MN} \delta[m] \delta[n], \quad (4.38)$$

with $0 \leq n \leq M_1 - 1$ and $0 \leq m \leq N_1 - 1$. The window function \tilde{g} is also periodic with period L , a fact that was assumed by Wexler and Raz in their derivation of the discrete-time Gabor expansion⁷. Equation (4.38) can be rewritten in matrix form as

$$\mathbf{G}_\gamma \vec{\gamma} = \begin{bmatrix} \frac{L}{MN} & 0 & \cdots & 0 \end{bmatrix}^T, \quad (4.39)$$

⁶Again, this is the exact opposite of the PF framework, where the overdetermined linear system representing the oversampled Gabor transform has *no* exact solutions – thus the motivation for a least-squared-error solution, which is provided by the OS-LSM-GT.

⁷The Wexler-Raz begins with a signal \tilde{x} with finite length L , constructs a pseudo-periodic signal \tilde{x} with period L , and then makes use of this periodicity to convert the continuous-time Gabor expansion into a discrete-time Gabor expansion. The same construction of a pseudo-periodic signal is used to convert the continuous-time Fourier transform into the DFT.

where the elements of \mathbf{G}_γ are given by

$$\mathbf{G}_\gamma(nN_1 + m, k) = \tilde{g}^*[k + nM]e^{\frac{-j2\pi mNk}{L}},$$

and where the tilde has been dropped from $\tilde{\gamma}$ to emphasize that $\tilde{\gamma}$ is a vector. In equation (4.39), \mathbf{G}_γ is $(M_1N_1 \times L)$ and $\tilde{\gamma}$ is $(L \times 1)$. Because $M_1N_1 < L$ in the WR framework for the oversampled case, the linear system (4.39) is again underdetermined and one must use some criterion to select a “best” solution out of the infinitely many exact solutions to (4.39). Once the “best” $\tilde{\gamma}$ is selected, the oversampled Gabor transform coefficients can be calculated using 4.36. This method of calculation, involving the selection of $\tilde{\gamma}$ to satisfy the discrete-time biorthogonality condition (4.39) and the calculation of C_{mn} using inner products between the signal \tilde{x} and time-frequency-shifted versions of $\tilde{\gamma}$, will be referred to here as the **periodized biorthogonal function method (PBFM)**.

Two methods have been suggested to select an “optimal” solution $\tilde{\gamma}_{opt}$ to (4.39). The first is to find the solution that minimizes the energy $E_\gamma = \tilde{\gamma}^H \tilde{\gamma}$ of the biorthogonal function vector $\tilde{\gamma}$. This solution is given by

$$\tilde{\gamma}_{opt} = \tilde{\gamma}_{min} = \mathbf{G}_\gamma^H (\mathbf{G}_\gamma \mathbf{G}_\gamma^H)^{-1} \tilde{v}_0, \quad (4.40)$$

where $\tilde{v}_0 = \left[\frac{L}{MN} \ 0 \ \dots \ 0 \right]^T$ is the constant vector on the right-hand side of equation (4.39). The matrix $\mathbf{G}_\gamma^H (\mathbf{G}_\gamma \mathbf{G}_\gamma^H)^{-1}$ is the *pseudoinverse* of \mathbf{G}_γ (the expression differs from the pseudoinverse defined in (3.16), because (4.39) is an underdetermined system, whereas (3.16) is an overdetermined system).

Qian, Chen, and Li [QCL92] showed that $\tilde{\gamma}_{min}$ is also the biorthogonal function that best approximates the Gabor expansion window function \tilde{g} (assuming \tilde{g} has unit energy), in the sense that

$$\tilde{\gamma}_{min} = \min_{\tilde{\gamma}: \mathbf{G}_\gamma \tilde{\gamma} = \tilde{v}_0} \left\| \frac{\tilde{\gamma}}{\|\tilde{\gamma}\|} - \tilde{g} \right\|^2.$$

Finding $\tilde{\gamma}_{min}$ that best approximates \tilde{g} is useful for the following reason: if $\tilde{\gamma}_{min} \approx k\tilde{g}$ for some constant k , then

$$\tilde{x} = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} C_{mn} \tilde{g}_{mn} \quad (4.41)$$

$$= \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \langle \tilde{x}, (\tilde{\gamma}_{min})_{mn} \rangle \tilde{g}_{mn} \quad (4.42)$$

$$\approx \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} k \langle \tilde{x}, \tilde{g}_{mn} \rangle \tilde{g}_{mn}, \quad (4.43)$$

and the Gabor expansion is an “orthogonal-like” representation [QC93].

The other method of finding an “optimal” solution to (4.39) is to use the extra degrees of freedom in the underdetermined system of (4.39) to impose additional constraints on $\vec{\gamma}_{opt}$, thereby shaping $\vec{\gamma}_{opt}$ according to some optimizing criteria. This method was explored by Wexler and Raz in [WR90], where they forced various components of $\vec{\gamma}_{opt}$ to zero. However, it is unclear what guidelines should be used to determine *which* components and *how many* components of $\vec{\gamma}_{opt}$ to force to zero. Consequently, this method appears rather arbitrary in the absence of clear and understandable criteria for optimizing the shape of $\vec{\gamma}_{opt}$.

4.6.2 Direct solutions

Whereas the majority of published work [QC92, QC93, QC94a, QCL92, WR90] has focused on the solution of (4.34) by means of the PBFM using (4.36), it is also interesting to consider direct solutions to (4.34). Recall that the WR framework represents the Gabor expansion

$$\vec{x} = \mathbf{G}_o \vec{c} \quad (4.44)$$

as an underdetermined system, with \mathbf{G}_o having dimensions $(L \times MN)$, and with $L = 0.25 * MN$ (assuming $\alpha = \beta = 0.5$). Note that \mathbf{G}_o in (4.44), the signal subspace matrix, is *not* the same matrix as \mathbf{G}_γ in (4.39).

Direct solution of (4.44) can be accomplished using the same two methods as described above for the solution of (4.39). First, the solution \vec{c}_{min} that minimizes the energy $E_{\vec{c}} = \vec{c}^H \vec{c}$ of the coefficient vector \vec{c} over all solutions of (4.44) can be found using the pseudoinverse of \mathbf{G}_o :

$$\vec{c}_{min} = \mathbf{G}_o^H (\mathbf{G}_o \mathbf{G}_o^H)^{-1} \vec{x}. \quad (4.45)$$

It should be noted that this type of solution was proposed by Daubechies [Dau91] in the form of a linear operator for the case of continuous-time signals.

Second, similar to the zero-forcing method of Wexler and Raz described above, one can force certain components of \vec{c}_{opt} to zero. Here, zero-forcing takes on a much clearer meaning than it does within the PBFM: if certain Gabor transform coefficients in \vec{c}_{opt} are known to be zero, this information takes the form of additional constraints in selecting the solution of (4.44). In fact, one can think of the Gabor expansion in the PF framework as the system (4.44), but with so many components of \vec{c} forced to zero that the system becomes overdetermined.




Figure 4.10: Minimum energy biorthogonal function $\gamma_{min}[k]$, as compared to scaled version of original one-sided exponential window $g[k]$ (with parameter $\lambda = 1$).

4.6.3 Examples

Three of the above methods for calculating the oversampled Gabor transform within the WR framework will now be applied to Signals 1 and 2 from sections 4.3–4.4. The parameters that will be used for the WR framework are as follows: $M = 64$, $N = 16$, $L = 256$, and $\Delta = \frac{1}{32}$.

First, consider the PBFM solution using the minimum energy biorthogonal function vector $\vec{\gamma}_{min}$. Figure 4.10 plots the magnitude of $\vec{\gamma}_{min}$ as calculated from (4.40) and compares this with a scaled version⁸ of the magnitude of the one-sided exponential window function \vec{g} (parameter $\lambda = 1$).

Figure 4.11 displays the magnitudes of the oversampled Gabor transform coefficients and the corresponding contour plots when the PBFM with $\vec{\gamma}_{min}$ is applied to Signals 1 and 2. This solution will be termed the “minimum biorthogonal function energy” (MBE) solution.

A few important conclusions can be drawn from an examination of (4.11):

- There are many more transform coefficients within the WR framework. Notice

⁸The scaling factor is 0.01 and is necessary because the original window function \vec{g} does not have unit energy in discrete time (\vec{g} is comprised of the samples of a window function with unit energy in *continuous* time).

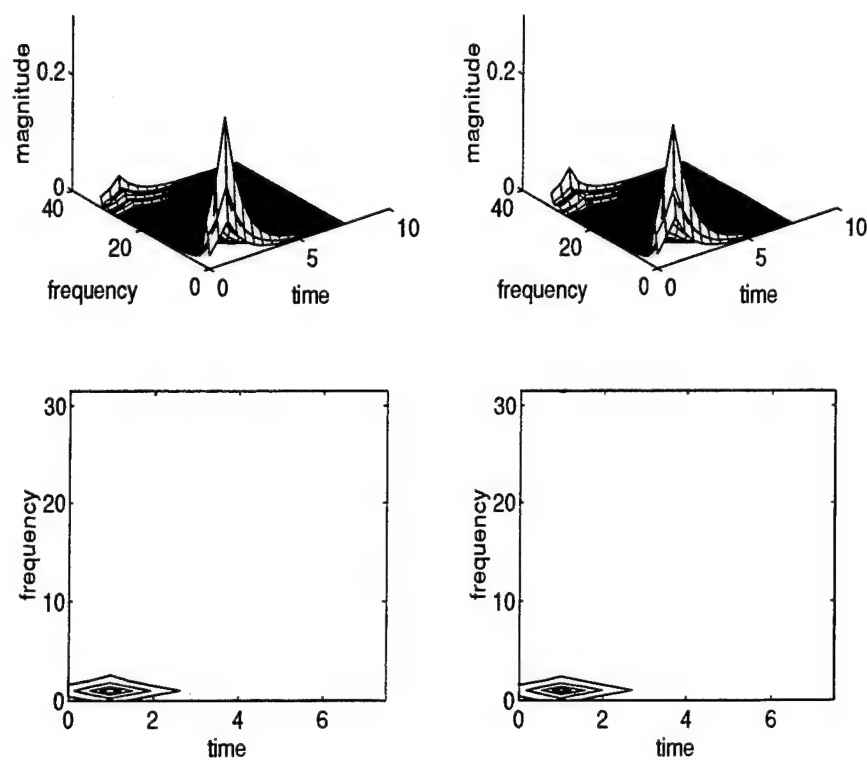


Figure 4.11: Left-hand side: Gabor transform coefficient magnitudes and corresponding contour plot for Signal 1, using PBFM with $\tilde{\gamma}_{min}$. Right-hand side: coefficient magnitudes and corresponding contour plot for Signal 2, using PBFM with $\tilde{\gamma}_{min}$.

that the frequencies of the coefficients extend to 31.5, and the times extend to 7.5. By comparison, the maximum time and maximum frequency for any of the coefficients in the PF framework was 3.5.

- Although the transform coefficient magnitudes do not exactly match that of the signal components (which have magnitude 1 for both Signal 1 and Signal 2), they are on the same order of magnitude. Moreover, the two plots are nearly identical for the two signals. Both of these observations indicate that the PBFM has good coefficient stability.
- The PBFM localizes the signal components for both signals to the correct (t, f) coordinates. However, just as with the TW-PCM and with the SSTFT, there is some blurring of the main peak.

Figure 4.12 displays the magnitudes of the oversampled Gabor transform coefficients and the corresponding contour plots when the minimum-coefficient-energy (MCE) solution is applied to Signals 1 and 2. The plots are very similar to those for Figure 4.11, and the same conclusions can be drawn.

Finally, zero-forcing of the transform coefficients (to be referred to as the ZF solution) is applied to Signals 1 and 2 in Figure 4.13. The ZF solution was obtained by forming the submatrix $(\mathbf{G}_o)_{zf}$ from the first 256 columns of \mathbf{G}_o and then solving the critically determined linear system

$$\vec{x} = (\mathbf{G}_o)_{zf}(\vec{c})_{zf}.$$

This is equivalent to forcing to zero all the transform coefficients corresponding to component arrival times greater than 1.5. Since the components in Signals 1 and 2 have arrival times around 1, this form of zero-forcing should not cause any inherent problems in the ZF solution.

Examination of Figure 4.13 reveals that the ZF solution still shows good coefficient stability, but there is a great deal of blurring in the time direction. However, the amount of blurring is exaggerated because of the shorter time axis; in reality, the blurring in Figure (4.13) is about the same as the blurring in Figures (4.11) and (4.12), perhaps even less in the frequency direction.

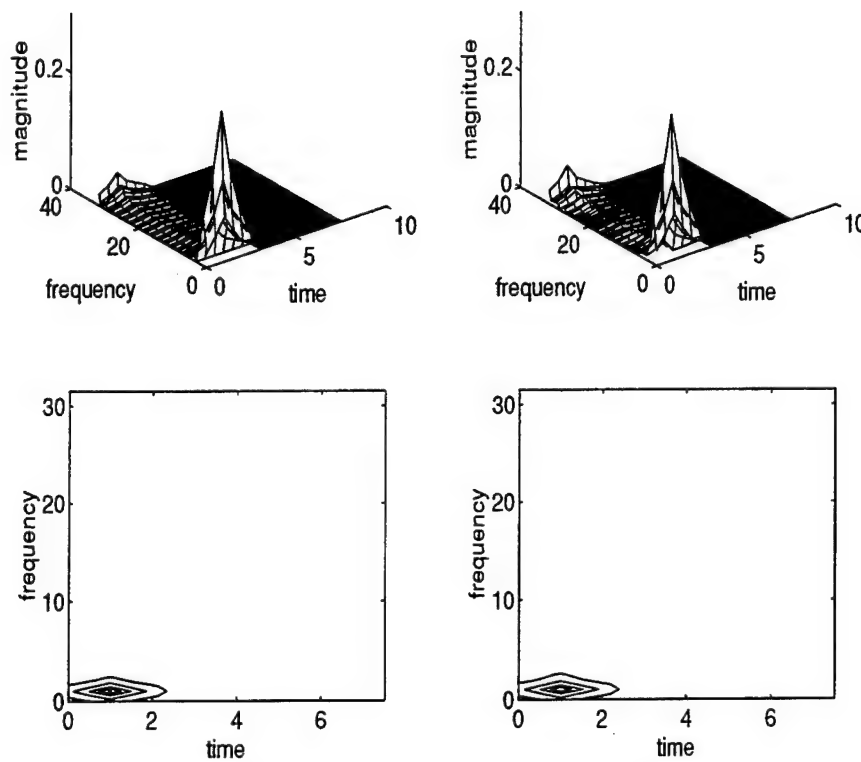


Figure 4.12: Left-hand side: Gabor transform coefficient magnitudes and corresponding contour plot for Signal 1, using minimum coefficient energy (MCE) solution. Right-hand side: coefficient magnitudes and corresponding contour plot for Signal 2, using MCE solution.

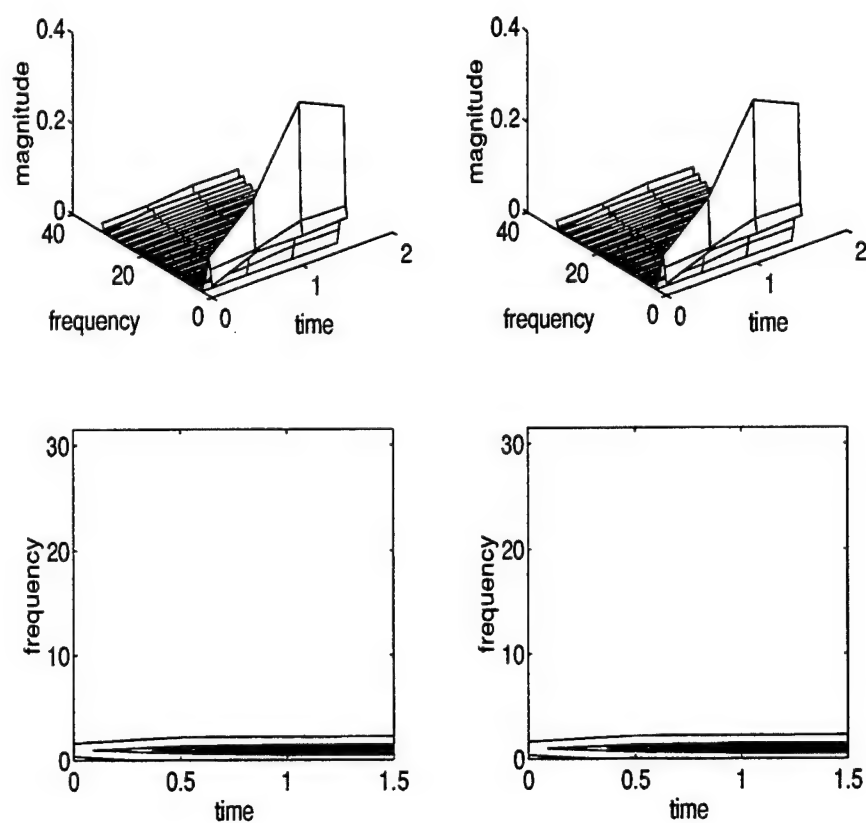


Figure 4.13: Left-hand side: Gabor transform coefficient magnitudes and corresponding contour plot for Signal 1, using zero-forced (ZF) solution. Right-hand side: coefficient magnitudes and corresponding contour plot for Signal 2, using ZF solution.

4.7 Analysis of the Oversampled Gabor Transform in the WR Framework

4.7.1 Stability analysis

The MBE, MCE, and ZF solutions for the oversampled Gabor transform within the WR framework can be compared to the five transforms listed in Table 4.1 for transform stability. The transform matrices corresponding to the respective solutions are as follows:

- For the MBE solution, the transform matrix is the biorthogonal function matrix Γ_W formed from time-frequency-shifted versions of the biorthogonal function $\vec{\gamma}_{min}$.
- For the MCE solution, the transform matrix is the pseudoinverse

$$\mathbf{B}_W = \mathbf{G}_o^H (\mathbf{G}_o \mathbf{G}_o^H)^{-1}.$$

- For the ZF solution, the transform matrix is the pseudoinverse

$$\mathbf{B}_{WZ} = (\mathbf{G}_o)_{zf}^H \left((\mathbf{G}_o)_{zf} (\mathbf{G}_o)_{zf}^H \right)^{-1}.$$

Recalling that the coefficient stability of the transform coefficient estimates (when the signal \vec{x} is embedded in complex, white Gaussian noise) is determined by the covariance matrix $2\mathbf{R}\mathbf{R}^H$ (for a generic transform matrix \mathbf{R}), Table 4.5 describes the “sizes” of the covariance matrices corresponding to the three WR framework solutions. Because the dimensions of the covariance matrices are $MN \times MN$, it was necessary to consider a set of WR framework parameters based on a lower sampling rate in order to generate Table 4.5: $M = 32$, $N = 16$, $L = 128$, $\Delta = \frac{1}{16}$. Note that the covariance matrices of Table 4.5 are still much larger in dimension than the covariance matrices of Table 4.1.

Examination of Table 4.5 reveals what was indicated by the examples of the previous section, that the WR framework solutions for the oversampled Gabor transform possess good coefficient stability.

4.7.2 Accuracy analysis

The MBE, MCE, and ZF solutions for the oversampled Gabor transform within the WR framework can also be compared to the transforms and modifications examined

Measure	Covariance Matrix		
	$2\Gamma_W\Gamma_W^H$	$2\mathbf{B}_W\mathbf{B}_W^H$	$2\mathbf{B}_{WZ}\mathbf{B}_{WZ}^H$
condition number	4.2260e18	4.5744e18	3.9068e18
2-norm	0.1574	0.0774	1.5066
1-norm	0.5278	0.1832	2.4294
∞ -norm	0.5278	0.1832	2.4294
Frobenius norm	1.3834	0.4124	3.1052
max(diag(\cdot))	0.0316	0.0160	0.3232
avg(diag(\cdot))	0.0302	0.0086	0.0238
min(diag(\cdot))	0.0200	0.0054	0.0116

Table 4.5: Stability of coefficient estimates, $M=32$, $N=16$, $L=128$, $\Delta = \frac{1}{16}$

in section 4.4, by means of the performance measures proposed in that section. In this subsection, the parameters of the WR framework are set back to $M = 64$, $N = 16$, $L = 256$, and $\Delta = \frac{1}{32}$.

Table 4.6 below lists the values of the three performance measures attained by the three WR framework solutions when applied to Signal 1.

Solution	Performance measure		
	NRE	CONC	CC
MBE	0.1415	0.6820	0.5109
MCE	ϵ	0.6877	0.5142
ZF	ϵ	0.8327	0.5994

Table 4.6: Comparison of performance measures for various solutions for the oversampled Gabor transform within the WR framework, as applied to Signal 1. Any value of NRE less than 10^{-9} is denoted by ϵ .

Examination of Table 4.6 reveals that all three solutions perform rather well for Signal 1 (the no-mismatch case), but not as well as the OS-LSM-GT. It should be noted that in the WR framework, all solutions should be *exact* solutions since the system is underdetermined. Thus, NRE should always be 0 regardless of the signal. That NRE is not 0 for the MBE solution is due to roundoff or computational error in the calculation of the inner products, not due to any theoretical deficiency of the MBE solution as compared to the other two solutions.

Table 4.7 below lists the values of the three performance measures attained by the three WR framework solutions when applied to Signal 2.

Solution	Performance measure		
	NRE	CONC	CC
MBE	0.1420	0.6563	0.4914
MCE	ϵ	0.6598	0.4945
ZF	ϵ	0.8060	0.5734

Table 4.7: Comparison of performance measures for various solutions for the over-sampled Gabor transform within the WR framework, as applied to Signal 2. Any value of NRE less than 10^{-9} is denoted by ϵ .

As indicated by the examples of the previous section, the performance measures for the three solutions do not degrade very much in the switch from Signal 1 to Signal 2.

Table 4.8 below lists the values of the three performance measures attained by the three WR framework solutions when applied to Signal 2 with complex Gaussian noise ($\sim \mathcal{N}(0, \mathbf{I})$) added in.

Solution	Performance measure		
	NRE	CONC	CC
MBE	0.1867	0.0701	0.2669
MCE	ϵ	0.0643	0.2522
ZF	0.7319	0.0948	0.1600

Table 4.8: Comparison of performance measures for various solutions for the over-sampled Gabor transform within the WR framework, as applied to Signal 2 with complex Gaussian noise $\sim \mathcal{N}(\vec{0}, \mathbf{I})$ added in. Any value of NRE less than 10^{-9} is denoted by ϵ .

Examination of Table 4.8 reveals that there is a dramatic drop in transform localization and accuracy (as measured by CONC and CC) for all three solutions when noise is added to Signal 2. Also, the NRE value for the ZF solution increases dramatically because a large portion of the Gaussian noise is orthogonal to the entire ZF signal subspace matrix $(\mathbf{G}_o)_{zf}$ (thus, the noise is not represented well by the ZF solution coefficients). Figure 4.14 illustrates the loss of accuracy in the MBE and MCE solutions when noise is added to Signal 2.

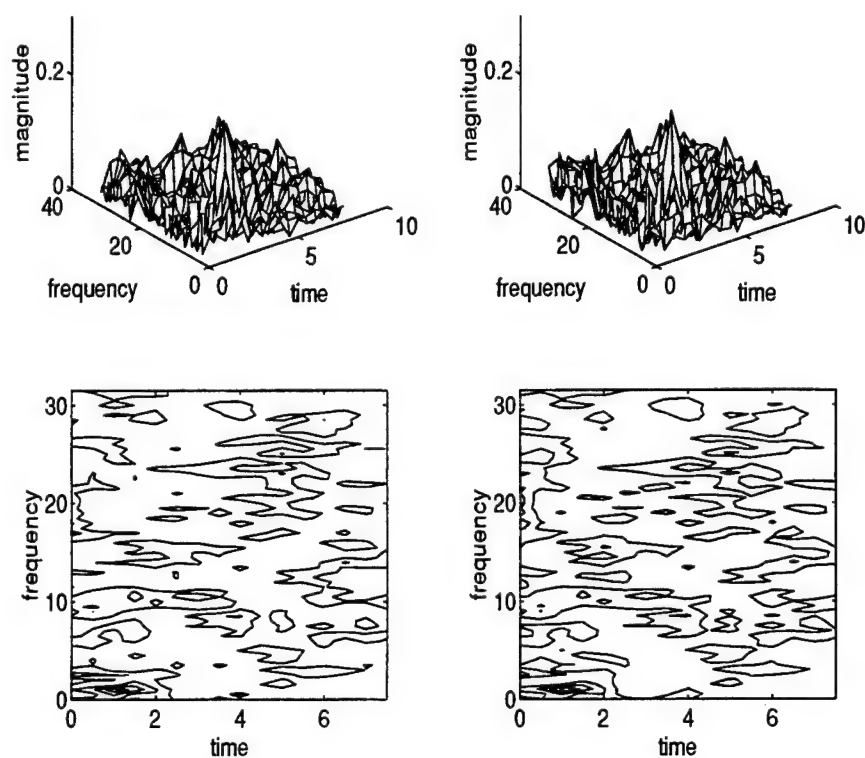


Figure 4.14: Left-hand side: Gabor transform coefficient magnitudes and corresponding contour plot for Signal 2 and noise, using MBE solution. Right-hand side: coefficient magnitudes and corresponding contour plot for Signal 2 and noise, using MCE solution.

While all the solutions maintain good coefficient stability, this drop in accuracy with the addition of noise is unsatisfactory and casts serious doubt on the use of the WR framework for transient signal detection problems.

4.8 Conclusions

This chapter analyzed several methods of improving upon the major weakness of the OS-LSM-GT, its lack of numerical stability. The first part of the chapter focused on several numerically-based modifications of the OS-LSM-GT while remaining within the PF framework. It was shown through examples and through performance analysis that none of these methods can simultaneously reconstruct the signal accurately and localize the time-frequency content of the signal accurately. If the former is most important, then the unmodified OS-LSM-GT is most effective; if the latter is most important, then either the OS-LSM-GT with transform whitening and PCM or the SSTFT should be used.

The second part of the chapter examined the calculation of the oversampled Gabor transform within another framework that has been the subject of much recent research, the Wexler-Raz (WR) framework. It was shown that calculation of the oversampled Gabor transform within the WR framework is most useful when one desires coefficient stability in the transform coefficients and accuracy in reconstructing the signal from the transform coefficients. However, there are several problems with using the WR framework to calculate the oversampled Gabor transform for transient signal detection problems:

1. The computational burden is much greater. For transient detection problems, it is especially disadvantageous to transform an L -length signal into MN coefficients if $MN > L$, since the purpose of the transform step is to simplify the problem. Yet, oversampling the Gabor transform within the WR framework dictates that $MN > L$.
2. The WR framework requires the calculation of many more Gabor transform coefficients than are necessary for cases where *a priori* knowledge of the signal's time-frequency characteristics can be applied. This was demonstrated in the example of section 4.5.
3. Calculating the oversampled Gabor transform within the WR framework requires the selection of a particular solution among an infinite set of exact solutions to an underdetermined system. It is not always clear what criteria should be used to make this selection.

4. None of the solution methods for the oversampled Gabor transform within the WR framework that were explored in this chapter provided accurate or localized transform coefficients when applied to a noisy signal.

The basic question posed by this chapter – “Is there a version of the oversampled Gabor transform (or of any oversampled linear transform) that provides both accurate reconstruction of the signal and an accurate, localized representation of the time-frequency content of the signal while maintaining coefficient stability?” – remains unanswered. All of the transforms and modifications investigated in this chapter – which encompass almost all of the methods for calculating the oversampled Gabor transform that can be found in the present literature – fall short in at least one of the three areas.

The above conclusions, then, appear fairly bleak. However, an adaptive version of the oversampled Gabor transform will be investigated in the next chapter, and this “adaptive Gabor transform” (based on the principles of *matching pursuit* [MZ92, MZ93, QCC92, QC94b]) will prove to be vastly superior to any of the transforms and modifications investigated in this chapter in terms of representing the time-frequency content of signals with good accuracy and good localization, while still maintaining fairly good reconstruction properties. Whether there are any non-adaptive forms of the oversampled Gabor transform that can provide all these properties simultaneously is a question – albeit a difficult one, as demonstrated by this chapter – for further consideration.

The Gabor Transform In Transient Signal Classification

5.1 Introduction

This chapter focuses on the problem of transient signal *classification*. In other words, the problem is not to determine *whether* a transient signal is present (or not), but rather to determine what *type* of transient signal is present *after* it has been determined that *some* kind of signal is indeed present.

Section 5.2 sets up a simple transient classification problem using the PF framework as a base model. Section 5.3 presents two GLRT-based classifiers for the problem of section 5.2 and analyzes the statistics of these two classifiers.

Section 5.4 investigates the *adaptive Gabor transform* (AGT) as an alternative to the non-adaptive linear TFRs of Chapters 3 and 4. The AGT is applied to the sample signals of Chapter 4 and the performance measures of Chapter 4 are calculated for the resulting transform coefficients.

Section 5.5 presents a comparison among the two GLRT-based classifiers and a classifier based on the AGT. Classifier performance is evaluated through Monte Carlo simulations. Using the results of the Monte Carlo simulations, the optimal detector-classifier combination within the PF framework is proposed.

Finally, section 5.6 summarizes the conclusions of the chapter and presents some ideas for future work.

5.2 Framework for Transient Signal Classification

The model for transient signal classification that will be used in this chapter is based on the PF framework, assuming unknown signal component locations. Recall that the transform step for this particular problem is given by

$$\vec{z} = \mathbf{B}_o \mathbf{G}_o \vec{c} + \vec{w}, \quad (5.1)$$

where \mathbf{G}_o is a signal subspace matrix comprised of oversampled Gabor logons and $\mathbf{B}_o = (\mathbf{G}_o^H \mathbf{G}_o)^{-1} \mathbf{G}_o^H$ is the pseudoinverse of \mathbf{G}_o . \mathbf{B}_o , once again, is the transform matrix corresponding to the oversampled least-squares-method Gabor transform (OS-LSM-GT).

Recall also that the hypothesis test for transient signal detection within the PF framework in the unknown location case is given by

$$\begin{aligned} \mathbf{H}_0 : \vec{c} &= \vec{0} \\ &vs. \\ \mathbf{H}_1 : \vec{c} &\neq \vec{0}. \end{aligned} \quad (5.2)$$

It will be *assumed* for most of this chapter that hypothesis \mathbf{H}_1 has been chosen by a transient signal detector. The problem of transient signal *classification* is then to determine what type of signal has been detected.

The classification problem that will be examined in this chapter is a simple one: determine whether the signal has high-frequency components or low-frequency components. Here, *high-frequency* and *low-frequency* signal components are defined in terms of the maximum possible component frequency $(M - 1)\beta$ as follows: a *low-frequency* component is a component whose frequency f satisfies $0 \leq f \leq \frac{(M-1)\beta}{2}$, while a *high-frequency* component is a component whose frequency f satisfies $\frac{(M-1)\beta}{2} < f \leq (M - 1)\beta$. For example, with the value of $M = 8$ that was used in Chapter 3, $0 \leq f \leq 1.75$ is the range for low-frequency components and $1.75 < f \leq 3.5$ is the range for high-frequency components.

Since the oversampled Gabor logons only have frequency components that are integer multiples of $\beta = 0.5$, the above definition indicates the following for signals $\vec{x} = \mathbf{G}_o \vec{c}$ that are expressible as linear combinations of the oversampled Gabor logons:

A *low-frequency signal* \vec{x}_L is a signal that can be expressed as a linear combination of low-frequency Gabor logons, i.e., those logons corresponding

to frequencies f such that $0 \leq f \leq \frac{(M-2)\beta}{2}$. A *high-frequency signal* \vec{x}_H is a signal that can be expressed as a linear combination of high-frequency Gabor logons, i.e., those logons corresponding to frequencies f such that $\frac{M\beta}{2} \leq f \leq (M-1)\beta$.

For example, using $M = 8$, a low-frequency signal \vec{x}_L can be expressed as a linear combination of Gabor logons whose frequencies f satisfy $0 \leq f \leq 1.5$, while a high-frequency signal \vec{x}_H can be expressed as a linear combination of Gabor logons whose frequencies f satisfy $2.0 \leq f \leq 3.5$.

In order to facilitate a simple representation of the classification problem, it is advantageous to partition the oversampled signal subspace matrix \mathbf{G}_o into submatrices \mathbf{G}_L and \mathbf{G}_H , whose columns consist of low-frequency and high-frequency Gabor logons, respectively (the 'o' subscript has been dropped for convenience in the submatrices). Similarly, the transform coefficient vector \vec{c} can be partitioned into \vec{c}_L and \vec{c}_H . The signal $\vec{x} = \mathbf{G}_o \vec{c}$ can then be rewritten as

$$\vec{x} = \begin{bmatrix} \mathbf{G}_L & \mathbf{G}_H \end{bmatrix} \begin{bmatrix} \vec{c}_L \\ \vec{c}_H \end{bmatrix}. \quad (5.3)$$

With the partition given by (5.3), the classification problem is given by the following binary hypothesis test:

$$\begin{aligned} \mathbf{H}_{1a} : \vec{c}_L \neq \vec{0} \text{ and } \vec{c}_H = \vec{0} \\ \text{vs.} \\ \mathbf{H}_{1b} : \vec{c}_L = \vec{0} \text{ and } \vec{c}_H \neq \vec{0}. \end{aligned} \quad (5.4)$$

Arbitrarily, \mathbf{H}_{1b} will be considered the null hypothesis and \mathbf{H}_{1a} will be considered the alternative hypothesis.

5.3 GLRT-Based Classifiers for Transient Signal Classification

5.3.1 GLRT classifier

The hypothesis test given by (5.4) contains a composite alternative hypothesis *and* a composite null hypothesis, so the logarithmic form of the GLRT (given the transformed observations \vec{z}) is given by

$$t_1 = \max_{\vec{c}} \{2 \log p_{1a}(\vec{z} | \vec{c})\} - \max_{\vec{c}} \{2 \log p_{1b}(\vec{z} | \vec{c})\}. \quad (5.5)$$

Given that $\vec{z} \sim \mathcal{N}(\vec{c}, 2\mathbf{B}_o\mathbf{B}_o^H)$ and letting

$$(2\mathbf{B}_o\mathbf{B}_o^H)^{-1} = \mathbf{\Lambda}^{-1} = \mathbf{\Pi} = \begin{bmatrix} \mathbf{\Pi}_{LL} & \mathbf{\Pi}_{LH} \\ \mathbf{\Pi}_{HL} & \mathbf{\Pi}_{HH} \end{bmatrix},$$

the term $2\log p_{1a}(\vec{z})$ in (5.5) can be written as

$$\begin{aligned} 2\log p_{1a}(\vec{z}) &= -MN \log 2\pi - \log |\mathbf{\Lambda}| \\ &\quad - \begin{bmatrix} \vec{z}_L^H - \vec{c}_L^H & \vec{z}_H^H \end{bmatrix} \begin{bmatrix} \mathbf{\Pi}_{LL} & \mathbf{\Pi}_{LH} \\ \mathbf{\Pi}_{HL} & \mathbf{\Pi}_{HH} \end{bmatrix} \begin{bmatrix} \vec{z}_L - \vec{c}_L \\ \vec{z}_H \end{bmatrix} \\ &= \kappa - (\vec{z}_L - \vec{c}_L)^H \mathbf{\Pi}_{LL} (\vec{z}_L - \vec{c}_L) + \vec{z}_H^H \mathbf{\Pi}_{HL} (\vec{z}_L - \vec{c}_L) \\ &\quad + (\vec{z}_L - \vec{c}_L)^H \mathbf{\Pi}_{LH} \vec{z}_H + \vec{z}_H^H \mathbf{\Pi}_{HH} \vec{z}_H \\ &= \kappa - (\vec{z}_L - \vec{c}_L + \mathbf{\Pi}_{LL}^{-1} \mathbf{\Pi}_{LH} \vec{z}_H)^H \mathbf{\Pi}_{LL} (\vec{z}_L - \vec{c}_L + \mathbf{\Pi}_{LL}^{-1} \mathbf{\Pi}_{LH} \vec{z}_H) \\ &\quad + \vec{z}_H^H (\mathbf{\Pi}_{HH} - \mathbf{\Pi}_{HL} \mathbf{\Pi}_{LL}^{-1} \mathbf{\Pi}_{LH}) \vec{z}_H, \end{aligned} \quad (5.6)$$

where $\kappa = -MN \log 2\pi - \log |\mathbf{\Lambda}|$ is constant. The last expression in (5.6) is maximized when

$$\vec{c}_L = \vec{z}_L + \mathbf{\Pi}_{LL}^{-1} \mathbf{\Pi}_{LH} \vec{z}_H,$$

which sets the second term equal to zero and results in a maximum value of

$$\begin{aligned} \max_{\vec{c}} \{2\log p_{1a}(\vec{z} | \vec{c})\} &= \kappa - \vec{z}_H^H (\mathbf{\Pi}_{HH} - \mathbf{\Pi}_{HL} \mathbf{\Pi}_{LL}^{-1} \mathbf{\Pi}_{LH}) \vec{z}_H \\ &= \kappa - \vec{z}_H^H \mathbf{\Lambda}_{HH}^{-1} \vec{z}_H, \end{aligned} \quad (5.7)$$

where the last equality stems from a well-known result regarding the inverse of a partitioned matrix [Sch91, for example].

Similar calculations to the above produce the expression for the second term in the log GLRT statistic:

$$\max_{\vec{c}} \{2\log p_{1b}(\vec{z} | \vec{c})\} = \kappa - \vec{z}_L^H \mathbf{\Lambda}_{LL}^{-1} \vec{z}_L. \quad (5.8)$$

Thus, the GLRT statistic for the classification problem of (5.4) is given by

$$t_1 = \vec{z}_L^H \mathbf{\Lambda}_{LL}^{-1} \vec{z}_L - \vec{z}_H^H \mathbf{\Lambda}_{HH}^{-1} \vec{z}_H. \quad (5.9)$$

To determine the statistics of t_1 , note that

$$\vec{z} = \begin{bmatrix} \vec{z}_L \\ \vec{z}_H \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \vec{c}_L \\ \vec{c}_H \end{bmatrix}, 2 \begin{bmatrix} \mathbf{\Lambda}_{LL} & \mathbf{\Lambda}_{LH} \\ \mathbf{\Lambda}_{HL} & \mathbf{\Lambda}_{HH} \end{bmatrix} \right), \quad (5.10)$$

so that

$$\vec{z}_L \sim \mathcal{N}(\vec{c}_L, 2\mathbf{\Lambda}_{LL}) \quad \text{and} \quad \vec{z}_H \sim \mathcal{N}(\vec{c}_H, 2\mathbf{\Lambda}_{HH}). \quad (5.11)$$

Given the statistical properties of \vec{z}_L , one can deduce that $\vec{z}_L^H \Lambda_{LL}^{-1} \vec{z}_L$ is noncentral χ^2 -distributed with MN degrees of freedom and noncentrality parameter $\nu_L = \vec{c}_L^H \Lambda_{LL}^{-1} \vec{c}_L$ under \mathbf{H}_{1a} and central χ^2 -distributed with MN degrees of freedom under \mathbf{H}_{1b} . Similarly, $\vec{z}_H^H \Lambda_{HH}^{-1} \vec{z}_H$ is noncentral χ^2 -distributed with MN degrees of freedom and noncentrality parameter $\nu_H = \vec{c}_H^H \Lambda_{HH}^{-1} \vec{c}_H$ under \mathbf{H}_{1b} and central χ^2 -distributed with MN degrees of freedom under \mathbf{H}_{1a} . Under either hypothesis, then, the statistic t_1 is the difference between a central χ^2 -distributed random variable and a noncentral χ^2 -distributed random variable. Note that the formation of the GLRT statistic t_1 involves the following steps:

1. Partition \vec{z} into \vec{z}_L and \vec{z}_H .
2. "Whiten" \vec{z}_L and \vec{z}_H separately, i.e., premultiply \vec{z}_L by $\Lambda_{LL}^{-\frac{1}{2}}$ and premultiply \vec{z}_H by $\Lambda_{HH}^{-\frac{1}{2}}$. Let $\vec{w}_L = \Lambda_{LL}^{-\frac{1}{2}} \vec{z}_L$ and $\vec{w}_H = \Lambda_{HH}^{-\frac{1}{2}} \vec{z}_H$ denote the whitened subvectors.
3. Calculate the energies $\vec{w}_L^H \vec{w}_L$ and $\vec{w}_H^H \vec{w}_H$ of the whitened subvectors.
4. Take the difference of the energies of the whitened subvectors to form $t_1 = \vec{w}_L^H \vec{w}_L - \vec{w}_H^H \vec{w}_H$.

Since t_1 is the difference of two χ^2 random variables, it is not possible to derive the probability distribution of t_1 in closed form [MP92]. However, it is possible to express the probability distribution of t_1 as an integral, using characteristic functions. This is done as follows:

1. Note that t_1 can be rewritten as $\vec{z}^H \mathbf{M} \vec{z}$, where

$$\mathbf{M} = \begin{bmatrix} \Lambda_{LL}^{-1} & \mathbf{0} \\ \mathbf{0} & \Lambda_{HH}^{-1} \end{bmatrix}.$$

2. Rewrite the expression for t_1 from step 1 in terms of $\vec{\zeta} = \sqrt{2} \Lambda^{-\frac{1}{2}} \vec{z}$:

$$t_1 = \vec{\zeta}^H \mathbf{M}_1 \vec{\zeta},$$

$$\text{where } \mathbf{M}_1 = \Lambda^{\frac{1}{2}} \mathbf{M} \Lambda^{\frac{1}{2}}.$$

3. Decompose \mathbf{M}_1 as $\mathbf{P} \mathbf{D} \mathbf{P}^{-1}$, where \mathbf{D} is diagonal and \mathbf{P} is unitary.
4. Rewrite the expression for t_1 from step 2 as

$$t_1 = \vec{\delta}^H \mathbf{D} \vec{\delta},$$

$$\text{where } \vec{\delta} = \mathbf{P}^{-1} \vec{\zeta} \sim \mathcal{N}(\mathbf{P}^{-1} \Lambda^{-\frac{1}{2}} \vec{c}, \mathbf{I}).$$

5. Rewrite the expression for t_1 from step 4 as

$$t_1 = d_1 \delta_1^2 + d_2 \delta_2^2 + \dots + d_{last} \delta_{last}^2,$$

where d_i is the i th diagonal element of \mathbf{D} and where $last = \frac{MN}{2}$. Note that this last expression for t_1 represents t_1 as a sum of independent random variables.

6. Calculate the characteristic function $\phi_{t_1}(v) = E\{e^{jvt_1}\}$ of t_1 as the product of the characteristic functions of the individual terms in the sum of step 5:

$$\phi_{t_1}(v) = \phi_{dd_1}(v) \times \phi_{dd_2}(v) \times \dots \times \phi_{dd_{last}}(v),$$

where $dd_i = d_i \delta_i$. Note that $\phi_{dd_i}(v) = \phi_{\delta_i}(d_i v)$ and that ϕ_{δ_i} is the characteristic function of a noncentral χ^2 random variable.

7. Determine the probability distribution $p(t_1)$ by taking the inverse Fourier transform of $\phi_{t_1}(v)$:

$$p(t_1) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \phi_{t_1}(v) e^{-jvt_1} dv. \quad (5.12)$$

Because the distribution of t_1 is dependent on the particular signals assumed in hypotheses H_{1a} and H_{1b} , it is necessary at this point to specify the parameters of the framework. Once again, the parameter values are assumed to be $M = N = 8$, $L = 256$, and $\Delta = \frac{1}{32}$ (see equation (3.2) to review the meanings of these parameters), and the window functions for the Gabor logons is a one-sided exponential window with parameter $\lambda = 1$. Under H_{1a} , \vec{c}_L has one nonzero component with magnitude 1, corresponding to the logon with (t, f) coordinates $(1, 1.5)$. Under H_{1b} , \vec{c}_H has one nonzero component with magnitude 1, corresponding to the logon with (t, f) coordinates $(1, 2)$. Assuming there is no location mismatch (i.e., t and f for all signal components are integer multiples of 0.5 and satisfy $0 \leq t \leq 0.5N$ and $0 \leq f \leq 0.5M$), the two signals described by \vec{c}_L and \vec{c}_H above represent the hardest possible pair of signals to identify correctly: one is low-frequency, one is high-frequency, and the two signals have components that are adjacent to each other on the oversampled Gabor transform grid. The time histories of the signals contain considerable overlap, since the signals have components with the same arrival time. The two signals are plotted in Figure 5.1.

Given the exact signals used for H_{1a} and H_{1b} , it is possible to plot the distributions for the GLRT statistic t_1 under H_{1a} and H_{1b} , using the seven-step method involving characteristic functions that was outlined above. The distributions are shown in Figure 5.2.




Figure 5.1: The real part of a low-frequency signal with (t, f) coordinates $(1, 1.5)$, plotted with the real part of a high-frequency signal with (t, f) coordinates $(1, 2)$. Both signals are based on a one-sided exponential window with parameter $\lambda = 1$.




Figure 5.2: Probability distributions of t_1 statistic under H_{1a} and H_{1b} , assuming signals as plotted in Figure 5.1.

It is clear from Figure 5.2 that the distributions of t_1 are nearly identical under \mathbf{H}_{1a} and \mathbf{H}_{1b} . Obviously, this indicates that the GLRT statistic t_1 will perform extremely poorly in classifying the two particular signals of Figure (5.1). It is well known [Fri90] that the GLRT statistic is not necessarily an optimal statistic for a given binary hypothesis test, and this is clearly one example.

5.3.2 Modified-GLRT classifier

Given the poor classification performance of the GLRT statistic t_1 , a modification of the GLRT statistic is proposed here. The modified-GLRT statistic t_2 is based on a partition of the whitened transform coefficients, $\vec{\zeta} = \Lambda^{-\frac{1}{2}} \vec{z}$:

$$t_2 = \vec{\zeta}_L^H \vec{\zeta}_L - \vec{\zeta}_H^H \vec{\zeta}_H. \quad (5.13)$$

Notice the difference in the formation of t_1 and t_2 :

- The GLRT statistic t_1 is formed by partitioning the transformed observations \vec{z} into \vec{z}_L and \vec{z}_H and then whitening the two subvectors *separately* into \vec{w}_L and \vec{w}_H . While \vec{w}_L and \vec{w}_H each have $2\mathbf{I}$ covariance matrices (i.e., the components within each subvector are independent), it is *not* the case that \vec{w}_L and \vec{w}_H are independent of each other.
- The modified GLRT statistic t_2 is formed by *first* whitening the *entire* transform vector \vec{z} into $\vec{\zeta}$, and *then* partitioning $\vec{\zeta}$ into $\vec{\zeta}_L$ and $\vec{\zeta}_H$. Thus, $\vec{\zeta}_L$ and $\vec{\zeta}_H$ again have $2\mathbf{I}$ covariance matrices, but they are *also* independent of each other.

Given the particular signals of Figure 5.1 for hypotheses \mathbf{H}_{1a} and \mathbf{H}_{1b} , the probability distributions for t_2 under \mathbf{H}_{1a} and \mathbf{H}_{1b} can be plotted using the same seven-step method involving characteristic functions that was detailed in the previous section. Here, however, steps 1–3 can be skipped since t_2 is already of the form $t_2 = \vec{\delta}^H \mathbf{D} \vec{\delta}$, with $\vec{\delta} = \vec{\zeta}$ and

$$\mathbf{D} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I} \end{bmatrix}.$$

The probability distributions for t_2 under \mathbf{H}_{1a} and \mathbf{H}_{1b} are plotted in Figure 5.3.

It is clear from Figure 5.3 that t_2 is an improvement over t_1 . Assuming a threshold of 0, the values of p_{fa} and p_d are both calculated (using integral approximations) to be approximately 0.182. However, the rates for t_2 , while much lower than the error rates for t_1 , are still not particularly low. This is a direct consequence of the




Figure 5.3: Probability distributions of t_2 statistic under H_{1a} and H_{1b} , assuming signals as plotted in Figure 5.1.

fact noted in Chapter 4, that the transform-whitened OS-LSM-GT, upon which t_2 is based, represents the time-frequency location of a one-component signal fairly accurately but not very locally (i.e., there is blurring involved).

In the next section, an entirely different approach to transient signal classification will be taken, based on an adaptive form of the Gabor transform.

5.4 Adaptive Gabor Transform (AGT)

Because of the problems encountered above by both the GLRT classifier (based on the OS-LSM-GT) and – to a lesser extent – the modified-GLRT classifier (based on the transform-whitened OS-LSM-GT), this section examines a classifier based on the *adaptive Gabor transform* (AGT), which calculates the Gabor transform in an entirely different manner than any of the methods mentioned in previous chapters.

5.4.1 Calculating the AGT

The calculation of the AGT, first described concurrently by Mallat and Zhang [MZ92, MZ93] and by Qian, Chen, and Chen [QCC92], is based on the idea of approximating vectors through *iterated* one-dimensional projections. Mallat and Zhang term this process “matching pursuit.”

The algorithm for calculating the AGT of a signal \vec{x} using the principles of matching pursuit is given as follows:

1. Let the set of oversampled Gabor logons comprising the columns of \mathbf{G}_o be denoted as $\mathcal{S} = \{\vec{g}_{mn} : 0 \leq m \leq 7, 0 \leq n \leq 7\}$, where the logons \vec{g}_{mn} are defined in equation (3.39) of Chapter 3.
2. Find the logon $\vec{g}_{mn} \in \mathcal{S}$ that has the largest¹ inner product with \vec{x} and denote this logon $\vec{g}^{(1)}$:

$$\vec{g}^{(1)} = \arg \left\{ \max_{m,n} |\langle \vec{x}, \vec{g}_{mn} \rangle| \right\}. \quad (5.14)$$

Since complex numbers are involved, the absolute value in (5.14) indicates a complex number magnitude. This step is equivalent to finding the maximum component of $\vec{z} = \mathbf{G}_o \vec{x}$.

¹*Largest* here means largest absolute value.

3. Calculate the *residual vector* $\vec{r}^{(1)} = \vec{x} - \langle \vec{x}, \vec{g}^{(1)} \rangle \vec{g}^{(1)}$. If the Gabor logons are normalized (i.e., $\langle \vec{g}_{mn}, \vec{g}_{mn} \rangle = 1$ for all m, n), then $\vec{r}^{(1)}$ is orthogonal to $\vec{g}^{(1)}$.
4. Repeat step 2, with $\vec{r}^{(1)}$ substituted for \vec{x} :

$$\vec{g}^{(2)} = \arg \left\{ \max_{m,n} \left| \langle \vec{r}^{(1)}, \vec{g}_{mn} \rangle \right| \right\}. \quad (5.15)$$

This is equivalent to finding the maximum component of the vector $\mathbf{G}_o \vec{r}^{(1)}$. Then calculate the residual vector $\vec{r}^{(2)} = \vec{r}^{(1)} - \langle \vec{r}^{(1)}, \vec{g}^{(2)} \rangle \vec{g}^{(2)}$.

5. At any iteration l , if it is the case that $\langle \vec{r}^{(l-1)}, \vec{g}_{mn} \rangle = 0 \quad \forall \vec{g}_{mn} \in S$ or that the residual vector $\vec{r}^{(l)}$ is calculated to be 0 (or 0 to within some tolerance ϵ), then **stop**. Otherwise, continue iterating for as many iterations as desired. At the l th iteration, the goal is to find the logon $\vec{g}^{(l)}$ such that

$$\vec{g}^{(l)} = \arg \left\{ \max_{m,n} \left| \langle \vec{r}^{(l-1)}, \vec{g}_{mn} \rangle \right| \right\}, \quad (5.16)$$

and then to calculate the l th residue vector $\vec{r}^{(l)} = \vec{r}^{(l-1)} - \langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle \vec{g}^{(l)}$. By definition, $\vec{r}^{(0)} = \vec{x}$.

6. After p iterations, the signal \vec{x} is given by the decomposition

$$\vec{x} = \sum_{l=1}^p \langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle \vec{g}^{(l)} + \vec{r}^{(p)} \quad (5.17)$$

The inner products $\langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle$ ($l = 1, 2, \dots, p$) form the AGT coefficients. If more than one of the inner products correspond to the same Gabor logon, then they are simply added together to form the AGT coefficient for that logon.

The two natural termination conditions mentioned in step 5 of the AGT algorithm correspond to slightly different situations:

- If, at the l th iteration, it is the case that $\langle \vec{r}^{(l-1)}, \vec{g}_{mn} \rangle = 0 \quad \forall \vec{g}_{mn} \in S$, then the residual vector $\vec{r}^{(l-1)}$ is orthogonal to every logon in S and the algorithm cannot reduce the squared error $R_{l-1} = \langle \vec{r}^{(l-1)}, \vec{r}^{(l-1)} \rangle$ any further.
- If, after the l th iteration, it is the case that $\langle \vec{r}^{(l)}, \vec{r}^{(l)} \rangle$ is calculated to be 0 (or 0 to within some tolerance ϵ), then the algorithm has reached its ultimate goal of decomposing the signal completely.

If the number of iterations is allowed to become arbitrarily large, then one of the two termination conditions will be reached. To show this, it is necessary to show first that the residual errors $R_l = \langle \vec{r}^{(l)}, \vec{r}^{(l)} \rangle$ form a monotonically decreasing sequence.

Proposition 5.1: *For any positive integer l , the residue vectors $\vec{r}^{(l-1)}$ and $\vec{r}^{(l)}$ have energies R_{l-1} and R_l , respectively, that satisfy $R_l \leq R_{l-1}$.*

Proof: Recall that $\vec{r}^{(l)}$ is defined as

$$\vec{r}^{(l)} = \vec{r}^{(l-1)} - \langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle \vec{g}^{(l)},$$

where $\vec{g}^{(l)}$ is defined as

$$\vec{g}^{(l)} = \arg \left\{ \max_{m,n} \langle \vec{r}^{(l-1)}, \vec{g}_{mn} \rangle \right\}.$$

It is either the case that $\langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle = 0$ or $\langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle \neq 0$. If it is the case that $\langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle = 0$, then $\vec{r}^{(l)} = \vec{r}^{(l-1)}$ and $R_{l-1} = R_l$. If $\langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle \neq 0$, then

$$\begin{aligned} R_l = \langle \vec{r}^{(l)}, \vec{r}^{(l)} \rangle &= \langle \vec{r}^{(l-1)}, \vec{r}^{(l-1)} \rangle - 2 \langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle^2 + \langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle^2 \\ &= \langle \vec{r}^{(l-1)}, \vec{r}^{(l-1)} \rangle - \langle \vec{r}^{(l-1)}, \vec{g}^{(l)} \rangle^2 \\ &\leq \langle \vec{r}^{(l-1)}, \vec{r}^{(l-1)} \rangle \\ &= R_{l-1}. \end{aligned}$$

■

Proposition 5.1 makes it possible to show that the sequence of squared residual errors R_l converges as $l \rightarrow \infty$.

Proposition 5.2: *The sequence $\{R_l\}_l$ of squared residual errors converges as $l \rightarrow \infty$.*

Proof: For $l = 1, 2, \dots$, the sequence $\{R_l\}_l$ of squared errors is a monotonically decreasing sequence (from Proposition 5.1), bounded above by $E = \vec{x}^H \vec{x}$ and bounded below by 0. It is well known [Rud76, theorem 3.14] that a bounded monotonic sequence converges. Thus, $\{R_l\}_l$ is a convergent sequence. ■

If the sequence $\{R_l\}_l$ converges to 0, then the second termination condition of the AGT algorithm is reached. If $\{R_l\}_l$ converges to a number greater than 0 (but less than or equal to $E = \vec{x}^H \vec{x}$), then the first termination condition of the AGT algorithm is reached. In the latter case, the limiting residual vector $\vec{r}^{(lim)}$ may still have an energy greater than the *minimum* squared error

$$SE_{min} = \langle \vec{x} - \mathbf{G}_o \mathbf{B}_o \vec{x}, \vec{x} - \mathbf{G}_o \mathbf{B}_o \vec{x} \rangle$$

resulting from the OS-LSM-GT. In other words, $\vec{r}^{(lim)}$ may be a vector that is orthogonal to the columns of \mathbf{G}_o but not equal to the error vector $\vec{x} - \mathbf{G}_o \mathbf{B}_o \vec{x}$ produced by the OS-LSM-GT ².

Finally, the following proposition covers the special case when the signal \vec{x} lies within $\text{span}\{\mathcal{S}\}$.

Proposition 5.3: *If the signal \vec{x} lies in within $\text{span}\{\mathcal{S}\}$ (i.e., $\vec{x} = \mathbf{G}_o \vec{c}$ for some \vec{c}), then the sequence $\{R_l\}_l$ converges to 0 and the AGT coefficients match the OS-LSM-GT coefficients.*

Proof: After every iteration, the residual vector can always be expressed as a linear combination of the Gabor logons; i.e.,

$$\vec{r}^{(l)} = \mathbf{G}_o \vec{c}_l,$$

where \vec{c}_l is some constant vector. Now the AGT algorithm only terminates when the remaining residual vector is orthogonal to every Gabor logon in \mathcal{S} (this holds true for either of the termination conditions). This orthogonality condition implies that

$$\mathbf{G}_o^H \vec{r}^{(lim)} = \mathbf{G}_o^H \mathbf{G}_o \vec{c}_{lim} = \vec{0}. \quad (5.18)$$

However, if $\vec{c}_{lim} \neq \vec{0}$, (5.18) implies that the Gram matrix $\mathbf{G}_o^H \mathbf{G}_o$ is singular, which is true if and only if the columns of \mathbf{G}_o are linearly dependent [Sch91, section 9.2]. Now the oversampled Gabor logons that comprise the set \mathcal{S} are linearly independent, so $\mathbf{G}_o^H \mathbf{G}_o$ is nonsingular. Therefore, $\vec{c}_{lim} = \vec{0}$ and the sequence $\{R_l\}_l$ converges to 0.

That the AGT coefficients converge to the OS-LSM-GT coefficients (which equal \vec{c} under the conditions of the proposition) is easily proved by contradiction. Suppose $\vec{c}_{lim} \neq \vec{c}$. Then, since the limiting residual vector is $\vec{0}$,

$$\begin{aligned} \vec{x} &= \mathbf{G}_o \vec{c}_{lim} = \mathbf{G}_o \vec{c} \\ \Rightarrow \mathbf{G}_o (\vec{c}_{lim} - \vec{c}) &= \vec{0}, \end{aligned}$$

²When \mathbf{G}_o has more rows than columns, as is the case with the PF framework, there are an infinite number of vectors \vec{v} such that $\mathbf{G}_o^H \vec{v} = 0$.

which implies that the columns of \mathbf{G}_o are linearly dependent, since $\vec{c}_{lim} - \vec{c} \neq \vec{0}$. As mentioned above, the columns of \mathbf{G}_o are linearly *independent*. Therefore, it must be the case that $\vec{c}_{lim} = \vec{c}$ and the AGT coefficients converge to the OS-LSM-GT coefficients under the conditions of the proposition. ■

5.4.2 Examples and performance measures for AGT

The AGT is now examined for the same signals and performance measures as used in Chapter 4. First, consider Signal 1, a one-component signal with one-sided exponential window (parameter $\lambda = 1$) located at (t, f) coordinates $(1, 1)$. Signal 1 is contained within $\text{span}\{S\}$, so Proposition 5.3 predicts that the AGT coefficients should converge to the actual coefficients of Signal 1. Figure 5.4 shows that the AGT coefficients do indeed converge to the actual coefficients of Signal 1 after just one iteration.




Figure 5.4: Magnitudes of AGT coefficients for Signal 1.

As shown in Table 5.1 below, the performance measures for the AGT (after one iteration) match those of the OS-LSM-GT for Signal 1 (both sets of measures take on ideal values):

Transform	Performance measure		
	NRE	CONC	CC
OS-LSM-GT	ϵ	1	1
AGT (1 iteration)	ϵ	1	1

Table 5.1: Performance measures for OS-LSM-GT and AGT as applied to Signal 1. This is the no mismatch case. Any value of NRE less than 10^{-9} is denoted by ϵ .

Next, consider Signal 2, a one-component signal with one-sided exponential window (parameter $\lambda = 1$) located at (t, f) coordinates $(1.05, 1)$. Recall that because of the instability of the OS-LSM-GT, the OS-LSM-GT coefficients for Signal 2 range in the thousands. Figure 5.5 shows the AGT coefficients of Signal 2 after 3 iterations and after 20 iterations, and it is clear that the AGT possesses good coefficient stability in the presence of information mismatch.

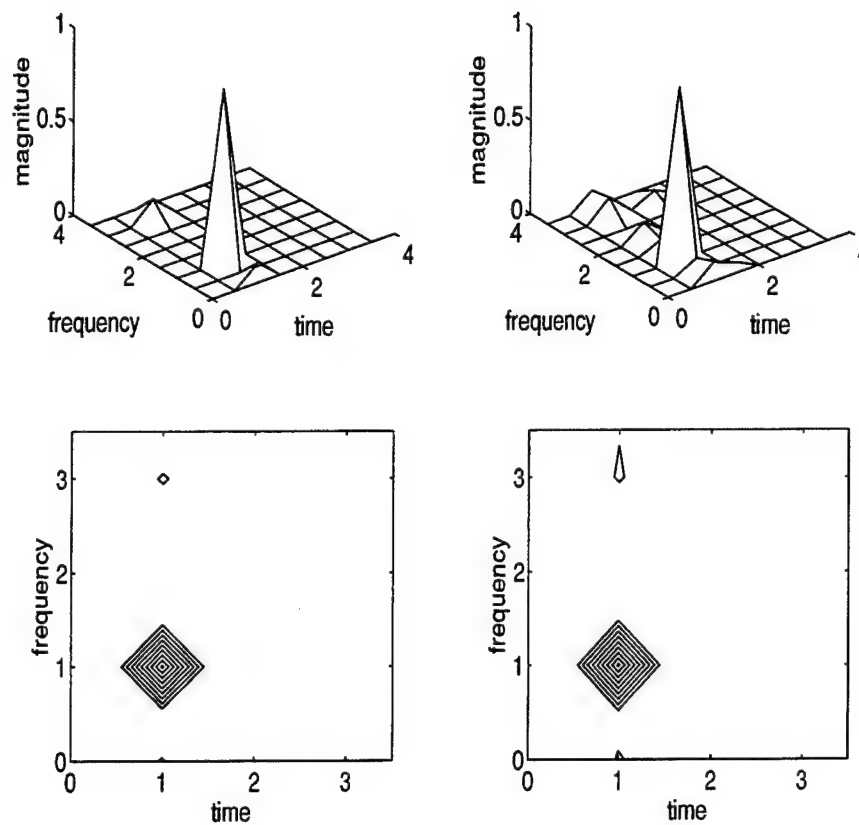


Figure 5.5: Left-hand side: magnitudes and corresponding contour plot of AGT coefficients for Signal 2, after 3 iterations. Right-hand side: magnitudes and corresponding contour plot of AGT coefficients for Signal 2, after 20 iterations.

Both cases show that the AGT coefficients represent the time-frequency content of the signal much more accurately than any of the non-iterative methods examined in Chapter 4. This is reflected in Table 5.2, which lists the corresponding performance measures for the AGT after 3 iterations, 20 iterations, and 100 iterations.

Transform	Performance measure		
	NRE	CONC	CC
OS-LSM-GT	0.0951	0.5027	0.2939
AGT (3 iterations)	0.3044	0.9726	0.9863
AGT (20 iterations)	0.2127	0.9467	0.9711
AGT (100 iterations)	0.1749	0.7724	0.8767

Table 5.2: Performance measures for OS-LSM-GT and AGT as applied to Signal 2.

Notice that as the number of iterations increases, NRE for the AGT decreases but CONC and CC also decrease. In other words, the increase in accuracy with respect to NRE resulting from increasing the number of iterations in the AGT algorithm is accompanied by a decrease in “localization accuracy” as measured by CONC and CC. Note, however, that the value of NRE attained by the AGT after even 3 iterations is perhaps acceptable given the excellent values the AGT attains for CONC and CC in that case.

Finally, consider Signal 2 with complex Gaussian noise ($\sim \mathcal{N}(\vec{0}, \mathbf{I})$) added in. Figure 5.6 shows the AGT coefficients of the noisy version of Signal 2, after 3 iterations and after 20 iterations.

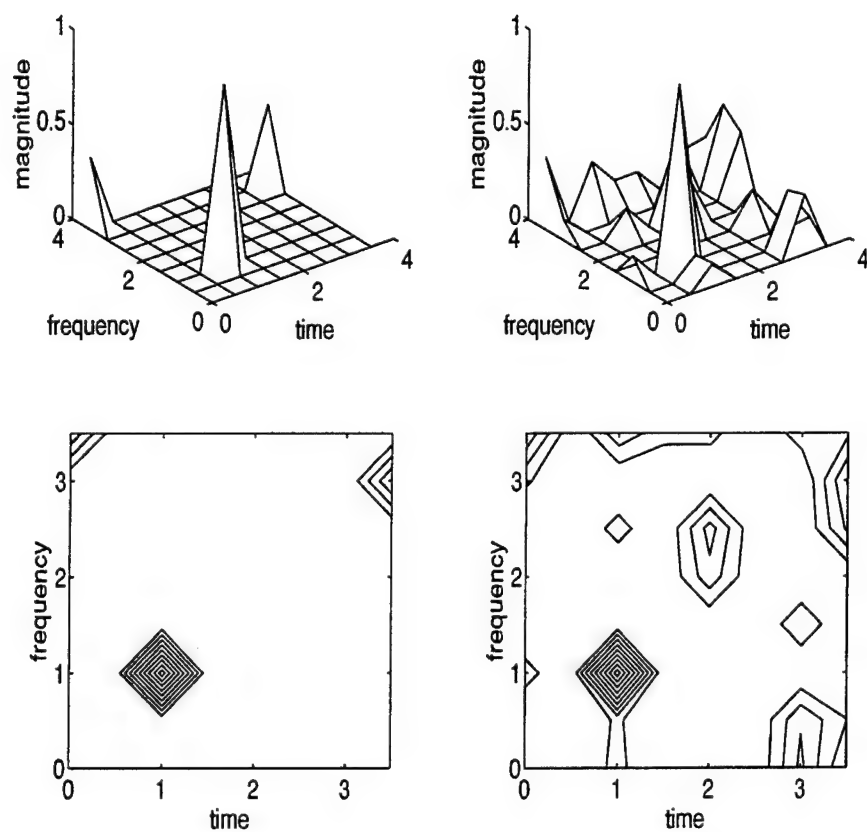


Figure 5.6: Left-hand side: magnitudes and corresponding contour plot of AGT coefficients (after 3 iterations) for Signal 2 added to complex noise $\sim \mathcal{N}(\vec{0}, \mathbf{I})$. Right-hand side: magnitudes and corresponding contour plot of AGT coefficients after 20 iterations for the same signal.

Examination of Figure 5.6 reveals that the localization of the signal component degrades considerably after 20 iterations in this case. This is confirmed in Table 5.3 below. Table 5.3 indicates that there is a considerable decrease in localization

Transform	Performance measure		
	NRE	CONC	CC
OS-LSM-GT	0.7628	0.3770	0.2650
AGT (3 iterations)	0.9270	0.7475	0.8635
AGT (20 iterations)	0.8785	0.4677	0.7025
AGT (100 iterations)	0.8571	0.3093	0.6014

Table 5.3: Performance measures for OS-LSM-GT and AGT as applied to Signal 2 added to complex noise $\sim \mathcal{N}(\vec{0}, \mathbf{I})$. The SNR is approximately -11.9062 dB.

accuracy with very little increase in reconstruction accuracy in return as the number of iterations in the AGT increases. Note that the NRE for the AGT after 3 iterations is already fairly close to the NRE for the OS-LSM-GT, the minimum NRE possible.

The above examples serve to illustrate that the AGT after a few iterations is vastly superior to any of the transforms and modifications examined in Chapter 4 in terms of representing the time-frequency content of signals with good accuracy and good localization; at the same time, the AGT achieves reasonable reconstruction accuracy ($NRE < 1$) after just a few iterations.

The main drawback to the AGT is that it calculates a localized representation (as it does for the signals above) *even if there is no signal*. For example, consider Figure 5.7, which plots the AGT coefficients after 3 iterations of the complex noise vector used in the last example above. Even though there is no signal present, the AGT coefficients of the noise seem to indicate that there are three signal components with appreciable magnitudes (ranging from 0.3 to 0.4 – recall that the magnitudes of the signal components in Signal 1 and Signal 2 are both 1). This property of the AGT precludes its use in signal *detection* schemes, since it does not distinguish “signal plus noise” from “noise alone” very well. However, once a signal has been detected, the AGT is very suitable for signal *classification* schemes, as will be demonstrated in the next section.




Figure 5.7: Magnitudes of AGT coefficients after 3 iterations for complex noise $\sim \mathcal{N}(\vec{0}, \mathbf{I})$.

5.5 Simulation Results

5.5.1 AGT classifier

The AGT coefficients as defined by the inner products in the decomposition of (5.17) can be used to form a classifier for the classification problem given by (5.4). Denoting the vector of AGT coefficients as \vec{a} and partitioning this vector into \vec{a}_L and \vec{a}_H in a manner similar to the partitioning of \vec{c} in (5.3), a logical AGT-based statistic for the hypothesis test in (5.4) is

$$t_3 = \vec{a}_L^H \vec{a}_L - \vec{a}_H^H \vec{a}_H. \quad (5.19)$$

The statistic t_3 simply calculates the energy of the low-frequency AGT coefficients and subtracts from this the energy of the high-frequency AGT coefficients. Because the AGT is an iterative algorithm, the statistics for t_3 are difficult if not impossible to determine analytically. Thus, the next section presents Monte Carlo simulations to compare the performance of t_3 to that of the GLRT-based classifiers t_1 and t_2 for the hypothesis test given by (5.4).

5.5.2 Monte Carlo simulations

This subsection presents the results of Monte Carlo simulations that compare the performance of the classification statistics t_1 , t_2 , and t_3 for the hypothesis test given by (5.4). Recall that t_1 represents the GLRT classifier, t_2 represents the modified GLRT (M-GLRT) classifier, and t_3 represents the AGT classifier. Again, all of the classifiers are based on the oversampled Gabor transform with parameters $\alpha = \beta = 0.5$. The AGT classifier is based on the AGT coefficients after three iterations.

In the first simulation, the low-frequency signal for hypothesis H_{1a} and the high-frequency signal for hypothesis H_{1b} were assumed to be the two signals plotted in (5.1). The thresholds for each of the detectors were set to 0. The 0 threshold does not represent a typical choice for a Neyman-Pearson decision rule. However, it is logical to assume that in the classification problem of (5.4), one type of error is not any more serious than the other — thus the choice of the 0 threshold, which equalizes the two error probabilities. For the statistic t_1 , the 0 threshold corresponds to an error probability of 0.493 for either hypothesis (see Figure 5.2). For the statistic t_2 , the 0 threshold corresponds to an error probability of 0.182 for either hypothesis (see Figure 5.3). As stated in the previous section, it is difficult if not impossible to determine the theoretical probability of error for the statistic t_3 .

The Monte Carlo simulations were then conducted as follows:

1. Generate the low-frequency signal shown in (5.1).
2. Add to the signal complex Gaussian noise $\vec{v} = \vec{v}_{re} + j\vec{v}_{im}$, with $\vec{v}_{re} \sim \mathcal{N}(\vec{0}, \mathbf{I})$ and $\vec{v}_{im} \sim \mathcal{N}(\vec{0}, \mathbf{I})$.
3. Calculate the statistics t_1 , t_2 , and t_3 . If any of them is less than zero, tally this as an error.
4. Go back to step 2 and repeat for a total of 5000 trials, keeping a running count of the errors generated by each statistic.
5. After all 5000 trials, divide the error counts by 5000 to obtain the sample miss probability associated with each statistic.
6. Generate the high-frequency signal shown in (5.1).
7. Add complex-Gaussian noise as in step 2 to the high-frequency signal.
8. Calculate the statistics t_1 , t_2 , and t_3 . If any of them is greater than or equal to zero, tally this as an error.

9. Go back to step 6 and repeat for a total of 5000 trials, keeping a running count of the errors generated by each statistic.
10. After all 5000 trials, divide the error counts by 5000 to obtain the sample false alarm probability associated with each statistic.

The results of the Monte Carlo simulation using the signals plotted in (5.1) are shown in Table 5.4 below.

Statistic	Error measure	
	P_m	P_{fa}
$t1$ (GLRT)	0.5036	0.4846
$t2$ (M-GLRT)	0.1836	0.1778
$t3$ (AGT)	0.0110	0.0142

Table 5.4: Monte Carlo error probabilities (5000 trials) for statistics $t1$ (GLRT), $t2$ (modified GLRT), and $t3$ (AGT, 3 iterations). The SNR is approximately -11.9062 dB. The (t, f) coordinates of the low-frequency signal are (1, 1.5); the coordinates of the high-frequency signal are (1, 2).

As expected, the classifier $t1$, based directly on the GLRT, is essentially useless for this problem. The classifier $t2$, based on the transformed-whitened OS-LSM-GT coefficients, performs acceptably. However, the classifier $t3$, based on the AGT coefficients after 3 iterations, performs exceptionally well, with less than a 2 percent error under either hypothesis. This exceptional performance was to be expected, given the analysis in the previous section concerning the AGT of Signal 1 (the no-mismatch case).

In the second Monte Carlo simulation, the signals were kept the same, except the frequency coordinates were varied to create a frequency mismatch. For the low-frequency signal (the alternative hypothesis, H_{1a}), the frequency was changed from 1.5 to 1.7; for the high-frequency signal (the null hypothesis, H_{1b}), the frequency was changed from 2.0 to 2.2. Again, the threshold was set to 0 for all three classifiers. The results of this simulation are shown in Table 5.5 below.

As expected, the simulation values of P_m increased greatly, as the higher frequency of the low-frequency signal increased the likelihood that the low-frequency signal was classified as a high-frequency signal. Conversely, the higher frequency of the high-frequency signal reduced the simulation values of P_{fa} for all three detectors. As in the first simulation example, the classifier based on the AGT coefficients

Statistic	Error measure	
	P_m	P_{fa}
$t1$ (GLRT)	0.5116	0.4788
$t2$ (M-GLRT)	0.4282	0.1050
$t3$ (AGT)	0.2608	0.0230

Table 5.5: Monte Carlo error probabilities (5000 trials) for statistics $t1$ (GLRT), $t2$ (modified GLRT), and $t3$ (AGT, 3 iterations). The SNR is approximately -9 dB. The (t, f) coordinates of the low-frequency signal are (1, 1.7); the coordinates of the high-frequency signal are (1, 2.2).

far outperformed the other two classifiers, an impressive result considering that the signals used for this second simulation example were frequency mismatched with the oversampled Gabor logons upon which all three classifiers were based.

In the third Monte Carlo simulation, the frequency of the high-frequency signal was changed from 2.2 to 1.8, making it almost the same as the frequency of the low frequency signal. This represents a much more difficult classification problem than either of the two examples above. The results of the Monte Carlo simulation for this problem, again using a threshold of 0 for all three classifiers, are shown in Table 5.6 below.

Statistic	Error measure	
	P_m	P_{fa}
$t1$ (GLRT)	0.5116	0.4882
$t2$ (M-GLRT)	0.4282	0.4052
$t3$ (AGT)	0.2608	0.2558

Table 5.6: Monte Carlo error probabilities (5000 trials) for statistics $t1$ (GLRT), $t2$ (modified GLRT), and $t3$ (AGT, 3 iterations). The SNR is approximately -9 dB. The (t, f) coordinates of the low-frequency signal are (1, 1.7); the coordinates of the high-frequency signal are (1, 1.8).

Note that the simulation values of P_m in Table 5.6 are the same as the values of P_m in Table 5.5, because the same low-frequency signal was used to determine P_m . However, with the frequency of the high-frequency signal *lowered* by 0.2 instead of raised by 0.2, the simulation values of P_{fa} increased to levels comparable to the simulation values of P_m . However, even with this very difficult classification problem,

the AGT-based classifier performs well.

It should be noted that, given the extremely good localization properties of the AGT noted in subsection 5.4.2, the performance of the AGT-based classifier can be improved by increasing the oversampling rate in the Gabor logons. For example, if the oversampled Gabor logons are separated by 0.1 instead of by 0.5, the AGT-based classifier would perform as well in the third simulation example above as in the first. Similarly increasing the oversampling rate for the modified-GLRT classifier (based on the transform-whitened OS-LSM-GT) would not be as effective because of the component blurring associated with transform whitening. One could propose to use the OS-LSM-GT directly with $\alpha = \beta = 0.1$ for the third simulation example above, but the instability of the OS-LSM-GT would be very evident if there were any mismatch. As revealed in subsection 5.4.2, the AGT does not suffer from instability in the presence of information mismatch, so increasing the oversampling rate will improve classifier performance without any harmful countereffects except that the computational load will be greater.

5.5.3 Optimal detector-classifier combination

The results of the Monte Carlo simulation in the previous subsection show that the AGT is clearly a superior means of signal classification. As noted previously, however, the AGT is not appropriate for use in signal detection problems, as the AGT coefficients do not distinguish “signal plus noise” from “noise alone” particularly well. This problem is not significant when the signal components have known locations in the time-frequency plane; in the KL case, the AGT can be used – likely with excellent results – in the same manner as any of the non-iterative methods of Chapter 3 were used, by focusing on the AGT coefficients at the known signal-component locations.

When the signal components have unknown locations, however, it is more important for a transform to distribute the energy of its coefficients when there is noise alone, and this is where the AGT falls short. Thus, the non-iterative methods of Chapter 3 are more appropriate in the UL case. The optimum detector-classifier combination in the UL case, then, would involve the following:

1. Use the UL detector of Chapter 3 to determine whether *any* signal is present. The threshold for this detection test should be set according to a Neyman-Pearson criterion, limiting P_{fa} .
2. If hypothesis \mathbf{H}_1 is chosen in Step 1, use the AGT-based classifier to determine what *type* of signal has been detected. The threshold for this test should be set

so that the error probabilities for either classification hypothesis are the same.

3. Denoting the latter error probability for the classification problem as P_e , the overall error probabilities are given by $1 - P_d$ if the signal is absent and by $(P_{fa}) + (P_d)(P_e)$ if a signal is present.

5.6 Conclusions

The chapter examined the problem of transient signal classification. The classification problem was defined within the PF framework as one in which the goal is to distinguish a low-frequency signal from a high-frequency signal. The GLRT statistic was derived for this classification problem, and it was found to be essentially useless. A modified form of the GLRT statistic, based on the transform-whitened OS-LSM-GT, was then proposed, and the modified-GLRT statistic was found to be more suitable to the classification problem (though not overwhelmingly so) than the statistic based on the GLRT alone.

Attention was then focused on the adaptive Gabor transform (AGT), which is based on the principles of matching pursuit. Using the performance measures and sample signals of Chapter 4, the performance of the AGT was found to be far superior to any of the non-iterative methods examined in that chapter in terms of representing the time-frequency content of signals with accuracy, localization, and coefficient stability. This superiority in the AGT was confirmed in Monte Carlo simulations, where the AGT-based classifier was found to be far superior to either of the two classifiers based on non-iterative transforms. It should be noted that the classification problem emphasizes the favorable properties of the AGT much more than the detection problem, as accuracy and localization of signal components is of primary importance. However, the AGT is not as suitable for the detection problem, where accuracy and localization of signal components in the transform is not as important as the ability to distinguish the transform representation of "signal plus noise" from the transform representation of "noise alone." Thus, a combination detector-classifier was proposed consisting of the unknown location (UL) detector based on the transform-whitened OS-LSM-GT proposed in the latter half of Chapter 3 and of the AGT-based classifier examined in this chapter.

Further research motivated by the work of this chapter should clearly focus on additional examination of the AGT. The effectiveness of the AGT in transient signal classification problems was only demonstrated in this chapter, not examined thoroughly. For example, it would be interesting to determine the optimal number of iterations that should be performed in the AGT for different classification problems.

It would also be interesting (though very difficult) to derive the statistical properties of the AGT coefficients in analytical form. Until this is done, the performance of the AGT can only be evaluated through simulations such as the Monte Carlo simulations used in this chapter. Also, although the AGT is not suitable for transient signal detection when component locations are unknown, it would be interesting to compare the performance of an AGT-based detector with the performance of the five detectors of Chapter 3 in the case where signal components are known; again, this type of comparison would entail numerical simulations such as Monte Carlo simulations, as the statistics of the AGT-based detector are not (at present) readily derived in analytical form.

Finally, the work of this chapter could be extended to the multiple-hypothesis case, where one desires to classify the signal among several frequency subbands instead of simply low-frequency and high-frequency subbands. It is likely that an AGT-based classifier would provide superior performance in the multiple-hypothesis case, for the same reasons that an AGT-based classifier provides superior performance in the two-hypothesis case.

Summary and Future Work

This dissertation examined the use of linear time-frequency representations – in particular, the Gabor transform and the short-time Fourier transform – for a certain class of transient signal detection and transient signal classification problems. The main conclusions of this study are detailed below, along with some suggestions for future work. Because ideas for future work were already discussed in the conclusion sections of Chapters 3–5, the discussion on future work here will focus on extending the results of this dissertation that involve the Gabor transform and the STFT to applications involving the third widely-used linear TFR, the wavelet transform.

Chapter 3 investigated the use of the oversampled Gabor transform in transient signal detection problems, using a model developed by Porat and Friedlander (the PF framework). The first part of the chapter focused on the case where the locations (in the time-frequency plane) of the signal components to be detected were assumed to be known. GLRT detectors based on the oversampled Gabor transform and on four critically sampled linear TFRs were compared, and it was found that the GLRT detector based on the oversampled Gabor transform was the most robust of all the detectors in the presence of every type of information mismatch (including frequency mismatch, arrival-time mismatch, and shape mismatch). The second part of the chapter focused on the case where signal component locations are assumed *unknown*. The GLRT detector based on the oversampled, least-squares method Gabor transform (OS-LSM-GT) for the “unknown location” (UL) case was derived, and it was compared to the corresponding GLRT detector based on the OS-LSM-GT for the “known location” (KL) case. It was found that the UL detector performed better than the KL detector unless the signal component locations were known to within adjacent coordinates on the oversampled Gabor lattice.

Friedlander and Porat in [FP92] investigated the use of the wavelet transform (WT) for transient signal detection within their framework. In this study, the signal

subspace matrix \mathbf{G} (or \mathbf{G}_o) was not comprised of Gabor logons (time-shifted and frequency-modulated versions of a basic function) but of wavelet “basis” functions (time-shifted and time-scaled versions of a basic function). The typical WT basis function was represented as

$$g_{mn}[k] = g_{mn}(k\Delta) = w(2^m k\Delta - n), \quad k = 0, 1, \dots, L-1, \quad (6.1)$$

with m (the scale variable) allowed to range from 0 to $M-1$ and n (the time variable) allowed to range from 0 to $N-1$. The function $w[k]$ represents the analyzing wavelet of the wavelet transform. Assuming then that the signal to be detected was a linear combination of wavelet functions of the form given in (6.1), Friedlander and Porat examined the performance of WT-based detectors in much the same way that they examined the performance of GT-based and STFT-based detectors in [PF92]. (Recall that the wavelet transform coefficients are calculated as inner products between the signal and the wavelet basis functions in (6.1). Thus, the transform matrix \mathbf{R} corresponding to the WT is simply the transpose of the signal subspace matrix formed using the functions (6.1) above.)

In [FP92], Friedlander and Porat investigated two types of analyzing wavelets, one based on the Daubechies orthonormal wavelet [Dau88] and another based on a slight variation of the one-sided exponential window function examined in Chapter 3. However, these two functions certainly do not represent all the analyzing wavelets that could be studied. For example, one frequently studied analyzing wavelet is the two-sided Gaussian window, given by

$$g_{mn}(t) = \sqrt{\frac{2\tau}{\pi}} e^{-\tau(t-n)^2 + j2\pi m(t-n)}. \quad (6.2)$$

Ironically, this was the original window function proposed by Gabor in [Gab46] while discussing the transform that now bears his name. In fact, the two-sided Gaussian window is often termed the “Gabor wavelet.” A typical Gabor wavelet is plotted in Figure 6.1. Although the Gabor wavelet does not have the abrupt onset and gradual decay typical of the shapes of many transient signals, it has been proposed as a model for certain types of transient phenomena in medical signal processing [Tut88, Tut89]. Thus, the performance of the Gabor wavelet within the PF framework merits further study. The simplicity and flexibility of the PF framework would allow such analysis to be conducted without too much variation from the analysis of Chapter 3.

Chapter 4 examined several methods of improving the major weakness of the oversampled least-squares-method Gabor transform, its lack of coefficient stability (i.e., the large variation in coefficient magnitudes with small changes in the signal or different realizations of noise). The first part of the chapter examined several




Figure 6.1: Two-sided Gaussian window $g_{mn}(t)$, with $m = 1$, $n = 0$.

methods for improving the coefficient stability of the OS-LSM-GT within the PF framework, including the widely used methods of rank reduction and regularization. Through examples and through a systematic performance analysis of some simple cases, it was found that none of these methods provided accurate, localized, *and* stable representations of the time-frequency content of the sample signals simultaneously. The transform-whitened OS-LSM-GT did prove to be stable, accurate (in terms of time-frequency representation), and fairly localized, but it had a high reconstruction error.

The second part of Chapter 4 examined the calculation of the oversampled Gabor transform within another framework that has been used a great deal in recent research, the Wexler-Raz (WR) framework. In the WR framework, the oversampled Gabor transform coefficients are found as the solution of an underdetermined system of equations. Thus, there are an infinite number of exact solutions, and four methods of selecting solutions were investigated in Chapter 4. With all four of the solution methods within the WR framework, it was found that the oversampled Gabor transform coefficients provided accurate, stable, and fairly localized time-frequency representations for signals without noise; however, the TFRs became very hard to interpret (i.e., dispersed instead of localized) when noise was added to the signals. The latter observation, along with a significant increase in the computational complexity of the oversampled Gabor transform within the WR framework over the corresponding computational complexity of the oversampled Gabor transform within the PF framework, prevents the effective use of the WR-framework version of the

oversampled Gabor transform in transient signal detection problems.

One adaptive method of improving the coefficient stability of the OS-LSM-GT in the presence of arrival-time mismatch involves the wavelet transform. The wavelet transform with the Haar basis [MCW93] can be used as a preprocessing tool to determine the arrival times of transient signal components having sharp onsets, such as those examined throughout this dissertation with one-sided exponential windows. In this procedure, the WT coefficients at the finest scale reveal the component arrival times very accurately, due to the discontinuities in the Haar functions. Once the component arrival times are determined, the Gabor logons can be adjusted accordingly (the Haar functions cannot be used directly, as they are rectangular functions that do not resemble the one-sided exponential window), so that there is no longer any arrival-time mismatch. This procedure, while it merits further investigation, does not change the inherent instability in the OS-LSM-GT, which still becomes apparent – even with no arrival-time mismatch – when noise is added to the signal.

Oversampling is hard to define for the wavelet transform, as the WT by definition has better time resolution at higher frequencies and better frequency resolution at lower frequencies. However, it is well known that the WT does suffer from coefficient instability in the presence of arrival-time mismatch [BE93]. Some research has been done in this area, but the topic is largely unexplored.

Chapter 5 investigated the use of the oversampled Gabor transform for transient signal classification problems defined within the PF framework. The specific classification problem was to distinguish high frequency signals from low frequency signals, once it was established that a signal of some kind was present. It was found for this specific problem that the GLRT statistic is essentially useless. A modified GLRT statistic based on the transform-whitened OS-LSM-GT was found to provide better – but not overwhelmingly superior – classification performance. Finally, the adaptive Gabor transform (AGT), based on the principles of matching pursuit, was defined, and a classifier based on the AGT was found to provide substantially better classification performance than either of the GLRT-based statistics.

The principles of matching pursuit were originally proposed by Mallat and Zhang [MZ92, MZ93] for use with the wavelet transform. Instead of Gabor logons, they defined a “dictionary” of time-frequency *atoms* for which to calculate the matching pursuit (MP) coefficients. The atoms were functions defined as

$$g_{m,n,a}(t) = \frac{1}{\sqrt{a}} g\left(\frac{t-n}{a}\right) e^{jmt}, \quad (6.3)$$

time-shifted, time-scaled, frequency-modulated versions of the window function $g(t)$. Note that the atoms of (6.3) simplify to the form of the (critically sampled) Gabor

logons when $a = 1$. The window function $g(t)$ was taken to be the two-sided Gaussian window (the Gabor wavelet) defined above. The main reason for this choice of $g(t)$ is that the Gaussian window provides the best joint time-frequency localization of any function [HBB92, Coh89, MZ92]. Thus, if nothing is known about the shapes of the signal components, then the two-sided Gaussian window is the best choice for $g(t)$.

While none of the original papers that proposed matching pursuit algorithms ([MZ92, MZ93, QC94b, QCC92]) investigated the use of such algorithms within the context of transient signal detection and classification, Frisch and Messer [FM91, FM92] examined transient signal detection using a related method. This latter method used the maximum wavelet transform coefficient as the detection statistic, which is equivalent to using the MP coefficients after one iteration. Through the use of Monte Carlo simulations, Frisch and Messer were able to demonstrate that their statistic provided better detection performance than the energy detector for their problem. As stated in the concluding section of Chapter 5, there are many interesting questions to be answered regarding the use of matching pursuit algorithms – with either Gabor logons or the more general time-frequency atoms of Mallat and Zhang – for transient signal detection and classification. One question of interest that was not mentioned in Chapter 5 is whether the two-sided Gaussian window truly provides the best performance (detection or classification) for matching pursuit algorithms when nothing is known about the waveform shapes of the signal components. A systematic study of this type has yet to be conducted.

In summary, this dissertation demonstrated that linear TFRs in general and the oversampled Gabor transform in particular are effective tools – when used with care – for transient signal detection and classification.

Bibliography

- [A⁺90] L. Auslander et al. A comparison of the Gabor and short-time Fourier transforms for signal detection and feature extraction in noisy environments. *Proceedings of SPIE Conference on Advanced Signal-Processing Algorithms, Architectures, and Implementations*, 1348:230–240, 1990.
- [Alt80] R. Altes. Detection, estimation, and classification with spectrograms. *Journal of the Acoustical Society of America*, 67(4):1232–1246, April 1980.
- [Bal92] R. Balart. Matrix reformulation of the Gabor transform. *Proceedings of SPIE Conference on Adaptive Signal Processing*, pages 447–457, 1992.
- [BE93] F. Bao and N. Erdol. On the discrete wavelet transform and shiftability. In *1993 27th Asilomar Conference*, pages 1442–1445, 1993.
- [BO92] R. Balart and R. Orr. Computational accuracy and stability issues for the finite, discrete Gabor transform. In *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pages 403–406, 1992.
- [BS90] B. Broder and S. Schwartz. Quickest detection procedures and transient signal detection. Technical Report 21, Princeton University, November 1990. Prepared for ONR.
- [CM80] T. Claasen and W. Mecklenbräuker. The Wigner distribution – a tool for time-frequency signal analysis, Parts I, II, and III. *Phillips Journal of Research*, 35:217–250, 276–300, 372–389, 1980.
- [CM84] T. Claasen and W. Mecklenbräuker. On the time-frequency discrimination of energy distributions: Can they look sharper than Heisenberg? In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, page 41.B.7, 1984.
- [Coh66] L. Cohen. Generalized phase-space distribution functions. *Journal of Mathematical Physics*, 7:781–786, 1966.

- [Coh89] L. Cohen. Time-frequency distributions – A review. *Proceedings of the IEEE*, 77(7):941–981, July 1989.
- [Coh92] L. Cohen. What is a multicomponent signal? In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages V.113–V.116, 1992.
- [CW89] H. Choi and W. Williams. Improved time-frequency representation of multicomponent signals using exponential kernels. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(6):862–871, June 1989.
- [Dau88] I. Daubechies. Orthonormal bases of compactly supported wavelets. *Communications in Pure and Applied Mathematics*, 41(7):909–996, 1988.
- [Dau90] I. Daubechies. The wavelet transform, time-frequency localization, and signal analysis. *IEEE Transactions on Information Theory*, 36(5):961–1005, September 1990.
- [Dau91] I. Daubechies. *Ten Lectures on Wavelets*. CBMS, SIAM, 1991.
- [EK91] T. Ebrahimi and M. Kunt. Image compression by Gabor expansion. *Optical Engineering*, 30(7):873–880, July 1991.
- [FM91] M. Frisch and H. Messer. Detection of a transient signal of unknown scaling and arrival time using the discrete wavelet transform. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages II.1313–II.1316, 1991.
- [FM92] M. Frisch and H. Messer. The use of the wavelet transform in the detection of an unknown transient signal. *IEEE Transactions on Information Theory*, 38(2):892–897, March 1992.
- [FP89] B. Friedlander and B. Porat. Detection of transient signals by the Gabor representation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(2):169–180, February 1989.
- [FP92] B. Friedlander and B. Porat. Performance analysis of transient detectors based on linear data transforms. *IEEE Transactions on Information Theory*, 38(2):665–673, March 1992.
- [FR94] S. Farkash and S. Raz. Linear systems in Gabor time-frequency space. *IEEE Transactions on Signal Processing*, 42(3):611–617, March 1994.
- [Fri90] B. Friedlander. Transient signal detection techniques. *U.S. Navy Journal of Underwater Acoustics*, 40(4):877–924, October 1990.

- [FZ] B. Friedlander and A. Zeira. Oversampled Gabor expansion into one-sided exponential functions. Submitted for publication.
- [Gab46] D. Gabor. Theory of communication. *Journal of the IEEE*, 93:429–457, November 1946.
- [GDL94] Z. Guo, L. Durand, and H. Lee. Comparison of time-frequency distribution techniques for analysis of simulated Doppler ultrasound signals of the femoral artery. *IEEE Transactions on Biomedical Engineering*, 41(4):332–342, April 1994.
- [GL83] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1983.
- [HBB92] F. Hlawatsch and G. Boudreaux-Bartels. Linear and quadratic time-frequency signal representations. *IEEE Signal Processing Magazine*, 9(2):21–68, April 1992.
- [Jan88] A. Janssen. The Zak transform: A signal transform for sampled time-continuous signals. *Phillips Journal of Research*, 43:23–69, 1988.
- [JP92] D. Jones and T. Parks. A resolution comparison of several time-frequency representations. *IEEE Transactions on Signal Processing*, 40(2):413–420, February 1992.
- [JW92] J. Jeong and W. Williams. Kernel design for reduced interference distributions. *IEEE Transactions on Signal Processing*, 40(2):402–412, February 1992.
- [KBB92] S. Kadambe and G. Boudreaux-Bartels. A comparison of the existence of “cross terms” in the Wigner distribution and the squared magnitude of the wavelet transform and the short time Fourier transform. *IEEE Transactions on Signal Processing*, 40(10):2498–2517, October 1992.
- [LHW94] J. Lu, D. Healy, and J. Weaver. Signal recovery and wavelet reproducing kernels. *IEEE Transactions on Signal Processing*, 42(7):1845–1848, July 1994.
- [LPA92] P. Loughlin, J. Pitton, and L. Atlas. Proper time-frequency energy distributions and the Heisenberg uncertainty principle. In *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pages 151–154, 1992.
- [Lue84] D. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, 2nd edition, 1984.

- [MCW93] F. Majid, R. Coifman, and M. V. Wickerhauser. *The Xwpl System Reference Manual: For version 1.3 of Xwpl*, December 1993.
- [MP92] A. Mathai and S. Provost. *Quadratic Forms in Random Variables*. Marcel Dekker, 1992.
- [MZ92] S. Mallat and Z. Zhang. Adaptive time-frequency decomposition with matching pursuits. In *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pages 7–10, 1992.
- [MZ93] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, December 1993.
- [NS82] A. Naylor and G. Sell. *Linear Operator Theory in Engineering and Science*. Springer-Verlag, 2nd edition, 1982.
- [Orr91] R. Orr. Computational assessment of Gabor representations. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 2217–2220, 1991.
- [Orr92] R. Orr. Finite discrete Gabor transform relations under periodization and sampling. In *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pages 395–398, 1992.
- [Orr93a] R. Orr. Derivation of the finite discrete Gabor transform by periodization and sampling. *Signal Processing*, 34(1):85–97, October 1993.
- [Orr93b] R. Orr. The order of computation for finite discrete Gabor transforms. *IEEE Transactions on Signal Processing*, 41(1):122–130, January 1993.
- [PF92] B. Porat and B. Friedlander. Performance analysis of a class of transient detection algorithms—a unified framework. *IEEE Transactions on Signal Processing*, 40(10):2536–2546, October 1992.
- [QC92] S. Qian and D. Chen. A general solution of biorthogonal analysis window functions for orthogonal-like discrete Gabor transform. In *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pages 387–389, 1992.
- [QC93] S. Qian and D. Chen. Discrete Gabor transform. *IEEE Transactions on Signal Processing*, 41(7):2429–2438, July 1993.

- [QC94a] S. Qian and D. Chen. Optimal biorthogonal analysis window function for discrete Gabor transform. *IEEE Transactions on Signal Processing*, 42(3):694–697, March 1994.
- [QC94b] S. Qian and D. Chen. Signal representation using adaptive normalized Gaussian functions. *Signal Processing*, 36(1):1–11, March 1994.
- [QCC92] S. Qian, D. Chen, and K. Chen. Signal approximation via data-adaptive normalized Gaussian functions and its applications for speech processing. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages I.141–I.144, 1992.
- [QCL92] S. Qian, K. Chen, and S. Li. Optimal biorthogonal functions for finite discrete-time Gabor expansion. *Signal Processing*, 27(2):177–185, May 1992.
- [Rud76] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 3rd edition, 1976.
- [RV91] O. Rioul and M. Vetterli. Wavelets and signal processing. *IEEE Signal Processing Magazine*, 8(4):14–38, October 1991.
- [Sch91] L. Scharf. *Statistical Signal Processing*. Addison-Wesley, 1991.
- [SF94] L. Scharf and B. Friedlander. Matched subspace detectors. *IEEE Transactions on Signal Processing*, 42(8):2146–2157, August 1994.
- [TK82] D. Tufts and R. Kumaresan. Estimation of frequencies of multiple sinusoids. *Proceedings of the IEEE*, 70(9):975–989, September 1982.
- [TKK82] D. Tufts, R. Kumaresan, and I. Kirsteins. Data adaptive signal estimation by singular value decomposition of a data matrix. *Proceedings of the IEEE*, 70(6):684–685, June 1982.
- [TS89] A. Thorpe and L. Scharf. Reduced rank methods for solving least squares problems. In *Proceedings of the Twenty-third Asilomar Conference on Signals, Systems, and Computers*, pages 609–613, 1989.
- [Tut88] F. Tuteur. Wavelet transformations in signal detection. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 1435–1438, 1988.
- [Tut89] F. Tuteur. Wavelet transformations in signal detection. In J. Combes, A. Grossmann, and P. Tchamitchian, editors, *Wavelets: Time-frequency Methods and Phase Space*, pages 132–138. Springer-Verlag, 1989.

- [WR90] J. Wexler and S. Raz. Discrete Gabor expansions. *Signal Processing*, 21:207–220, 1990.
- [ZAM90] Y. Zhao, L. Atlas, and R. Marks. The use of cone-shaped kernels for generalized time-frequency representations of nonstationary signals. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(7):1084–1091, July 1990.